

UK Regional Nowcasting using a Mixed Frequency Vector Autoregressive Model

Gary Koop^{1,2}, Stuart McIntyre^{1,2} and James Mitchell^{1,3}

¹Economic Statistics Centre of Excellence

²University of Strathclyde

³University of Warwick

ESCoE Discussion Paper 2018-07

June 2018

About the Economic Statistics Centre of Excellence (ESCoE)

The Economic Statistics Centre of Excellence provides research that addresses the challenges of measuring the modern economy, as recommended by Professor Sir Charles Bean in his Independent Review of UK Economics Statistics. ESCoE is an independent research centre sponsored by the Office for National Statistics (ONS). Key areas of investigation include: National Accounts and Beyond GDP, Productivity and the Modern economy, Regional and Labour Market statistics.

ESCoE is made up of a consortium of leading institutions led by the National Institute of Economic and Social Research (NIESR) with King's College London, innovation foundation Nesta, University of Cambridge, Warwick Business School (University of Warwick) and Strathclyde Business School.

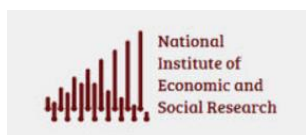
ESCoE Discussion Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the ESCoE, its partner institutions or the ONS.

For more information on ESCoE see www.escoe.ac.uk.

Contact Details

Economic Statistics Centre of Excellence
National Institute of Economic and Social Research
2 Dean Trench St
London SW1P 3HE
United Kingdom

T: +44 (0)20 7222 7665
E: escoeinfo@niesr.ac.uk



UK Regional Nowcasting using a Mixed Frequency Vector Autoregressive Model

Gary Koop^{1,2}, Stuart McIntyre^{1,2} and James Mitchell^{1,3}

¹Economic Statistics Centre of Excellence, ²University of Strathclyde, ³University of Warwick

Abstract

Data on Gross Value Added (GVA) are currently only available at the annual frequency for the UK regions and are released with significant delay. Regional policymakers would benefit from more frequent and timely data. The goal of this paper is to provide these. We use a mixed frequency Vector Autoregression (VAR) to provide, each quarter, nowcasts (i.e. forecasts of current GVA which is as yet unknown due to release delays) of annual GVA growth for the UK regions. The information we use to update our regional nowcasts comes from GVA growth for the UK as a whole as this is released in a more timely and frequent (quarterly) fashion. To improve our nowcasts we use entropic tilting methods to exploit the restriction that UK GVA growth is a weighted average of GVA growth for the UK regions. In this paper, we develop the econometric methodology and test it in the context of a real time nowcasting exercise.

Key words: Regional Growth; Nowcasting; Mixed frequency

JEL classification: C22, C52, C53, E01, R1

Contact Details

Stuart McIntyre
Fraser of Allander Institute
Department of Economics
University of Strathclyde
Glasgow G4 0QU
United Kingdom

Email: gary.koop@strath.ac.uk, s.mcintyre@strath.ac.uk, james.mitchell@wbs.ac.uk

This ESCoE paper was first published in June 2018.

© Gary Koop, Stuart McIntyre and James Mitchell

UK Regional Nowcasting using a Mixed Frequency Vector Autoregressive Model*

Gary Koop[†], Stuart McIntyre[‡] and James Mitchell[§]

Abstract: Data on Gross Value Added (GVA) are currently only available at the annual frequency for the UK regions and are released with significant delay. Regional policymakers would benefit from more frequent and timely data. The goal of this paper is to provide these. We use a mixed frequency Vector Autoregression (VAR) to provide, each quarter, nowcasts (i.e. forecasts of current GVA which is as yet unknown due to release delays) of annual GVA growth for the UK regions. The information we use to update our regional nowcasts comes from GVA growth for the UK as a whole as this is released in a more timely and frequent (quarterly) fashion. To improve our nowcasts we use entropic tilting methods to exploit the restriction that UK GVA growth is a weighted average of GVA growth for the UK regions. In this paper, we develop the econometric methodology and test it in the context of a real time nowcasting exercise.

*Thanks to ESCoE for financial support; and to our ESCoE referees and Aubrey Poon for helpful comments on an earlier draft of this paper. Thanks to Jeffrey Darko and Trevor Fenton, at the ONS, for helping us retrieve historical ONS data.

[†]Rimini Centre for Economic Analysis; Fraser of Allander Institute, Department of Economics, University of Strathclyde; Economic Statistics Centre of Excellence (gary.koop@strath.ac.uk)

[‡]Fraser of Allander Institute, Department of Economics, University of Strathclyde; Economic Statistics Centre of Excellence (s.mcintyre@strath.ac.uk)

[§]Warwick Business School, University of Warwick; Economic Statistics Centre of Excellence (James.Mitchell@wbs.ac.uk)

1 Introduction

The fact that data for many key macroeconomic variables are released only monthly, quarterly or annually, and even then they are released with delay, sparks interest in nowcasting. Our focus is on nowcasting nominal Gross Value Added (GVA) for the regions of the UK. GVA measures the increase in the value of the economy due to the production of goods and services.¹ For the UK as a whole, GVA data are released quarterly by the Office for National Statistics (ONS) with a first estimate currently released roughly two months after the end of the calendar quarter (along with the ONS’s “second estimate” of GDP). But regional GVA data are available only on an annual basis, with the initial release for a particular year currently occurring more than eleven months after the end of the year. Thus we have a frequency mismatch in terms of the data available at the national and regional levels. The advantage of incorporating high frequency (UK) data when nowcasting low frequency (regional) GVA data is clear: nowcasts can be produced every quarter (as new information about higher frequency variables such as UK GVA is released); and policymakers and decision makers do not have to wait until the end of the year to receive updated estimates of regional GVA.

To provide such estimates, we draw on and extend the growing literature on mixed frequency Vector Autoregressions (VARs). There are two main timing conventions used in this literature. The first, often called the stacked VAR approach, writes the VAR at the low frequency and the high frequency variables appear multiple times in each period (i.e. if there are R regions, then the dependent variables in the VAR will include the R regional variables plus four UK quarterly values). Pioneering stacked VAR papers include Ghysels (2016) and McCracken, Owyang and Sekhposyan (2016). The second, which can be called the state space VAR approach, writes the VAR at the high frequency and state space methods are used to fill in the missing observations (see, e.g., Eraker, Chiu, Foerster, Kim and Seoane, 2015, Schorfheide and Song, 2015 and Brave, Butters and Justiniano, 2016).

Our paper uses a stacked VAR. However, our approach deviates from conventional stacked VAR approaches in some ways and this requires the extension and adaptation of existing methods. A conventional stacked VAR approach exploits the information in many high frequency variables to update a single low frequency variable of interest (e.g. using many monthly macroeconomic variables to nowcast quarterly GDP growth).² In our re-

¹GVA plus taxes (less subsidies) on products is equivalent to gross domestic product (GDP).

²An exception is Ghysels, Grigoris and Ozkan (2017) which forecasts annual government expenditures and revenues for 48 US states using quarterly and monthly predictors.

gional nowcasting application, the frequency mismatch is reversed: we have many low frequency variables to nowcast (i.e. GVA growth for R UK regions) and a single high frequency predictor (i.e. quarterly UK GVA growth).³ This raises empirical challenges and motivates our use of the stacked VAR. That is, if we had used the state space VAR when our data set had so many low frequency variables, the number of missing observations would be large and estimation may be difficult. However, working with the stacked VAR means we are working with a higher dimensional VAR; this too imposes empirical challenges.

To overcome these challenges and add some empirically useful features, we extend standard stacked VAR methods in four ways. First, each quarter, as new releases of UK GVA data are produced, we entropically tilt towards these new releases so as to produce nowcasts of regional GVA which reflect this information. Secondly, we also use entropic tilting methods to exploit the fact that GVA growth for the UK as a whole should be (approximately, discussed further below) equal to a weighted average of regional GVA growth rates. Thirdly, we use Bayesian methods which allow for prior shrinkage (with the degree of shrinkage estimated from the data) so as to avoid over-parameterisation problems in our relatively large VAR. Finally, in many macroeconomic applications, there is strong evidence of volatility changes and, thus, conventional VAR applications often add multivariate stochastic volatility (e.g. Clark, 2011). Volatility issues have not been extensively considered in the mixed frequency VAR literature, but our specification allows for it.

Another contribution of this paper lies in the construction of a long time series of annual regional GVA data from 1966 to 2016 for the UK. The current regional GVA dataset from the ONS only begins in 1997. Details of how we combine these data with earlier sources are provided in the Data Appendix. Aware of data revisions, our ambition in putting together the database was to use, as close as possible to (over our out-of-sample window), first-release estimates of regional GVA and match these with the appropriate, similarly dated, data release for UK GVA. This means that in producing our nowcasts

Another exception is Mandalinci (2015) which is a regional UK GVA application with a similar frequency mis-match to ours. Both these papers use different econometric methods than those we employ and have a different empirical focus.

³Our methods could be extended to consider further high frequency predictors. Our focus in this paper, however, is on using UK GVA growth as a predictor. Importantly regional GVA adds up to UK GVA and we exploit this fact in our econometric specification. Other high frequency predictors that we might consider for nowcasting regional GVA, such as qualitative survey or labour market data, would not share this characteristic and therefore their utility in nowcasting regional GVA is essentially an empirical matter.

we are estimating our models on (as close as possible to, as explained in the Data Appendix) first-release estimates and evaluating each nowcast relative to the ONS’s first estimate of regional GVA. Clements and Galvao (2013) have advocated a similar use of ‘lightly revised’ data instead of using data from the latest-available (real-time) vintage.

Using these data and the stacked VAR, we carry out a real-time nowcasting exercise. At the beginning of each year, we provide unconditional (with respect to current year information) density forecasts of regional GVA growth for each region for the current year. These forecasts do, however, condition on data from previous years; and to acknowledge the publication lags of the regional data they are in effect two-year ahead forecasts, rather than just one year ahead, until late in the current year when the previous year’s regional data are published. Then, as each quarter of the current year passes by, and new UK-wide GVA data are released, we produce nowcasts of regional GVA growth which update the unconditional forecasts using entropic tilting methods. We find that these updated nowcasts are much more accurate than the initial unconditional forecasts, in terms of anticipating the ONS’s subsequent first releases for regional GVA growth. This provides evidence that the methods developed in this paper can be used to produce quarterly ‘flash’ (i.e. pre ONS first release) estimates of regional GVA growth where currently only annual estimates are available. They let us allocate national growth among the regions of the UK as soon as the quarterly UK figures are published, enabling the production of much more timely estimates of regional GVA growth. For instance, at the end of May 2017 we could already produce a nowcast of regional GVA growth for 2017, conditioning on 2017Q1 UK GVA data. The actual initial release of 2017 regional GVA by the ONS will not be until mid December 2018. In the time between May 2017 and December 2018, our nowcasts might be found useful by a regional policy-maker in giving an early and reliable signal of the state of the economy in their region.

2 The Econometrics of Regional Nowcasting

Our goal is to build an econometric model for nowcasting regional GVA growth using mixed frequency data which have annual observations for the regions and quarterly observations for the UK as a whole. Our nowcasts will be of annual growth rates, but they will be updated quarterly using entropic tilting methods. In this section, we first describe the stacked VAR we use. Next we describe the prior used to achieve prior shrinkage and avoid over-parameterisation concerns. Subsequently, we describe predictive and

posterior inference in the model. Finally, we describe how we implement the entropic tilting.

2.1 The Stacked VAR

First, we define our notation:

- $r = 1, \dots, R$ is an index for the UK regions.
- $t = 1, \dots, T$ is an index for time at the *annual* frequency.
- $Y_t^{r,A}$ is annual GVA for region r .
- $y_t^{r,A} = \left(\frac{Y_t^{r,A} - Y_{t-1}^{r,A}}{Y_{t-1}^{r,A}} \right)$ is annual GVA growth in region r .
- $Y_{t,q}^{UK}$ is UK GVA in the q^{th} quarter of year t where $q = 1, \dots, 4$.
- $y_{t,q}^{UK} = \left(\frac{Y_{t,q}^{UK} - Y_{t-1,q}^{UK}}{Y_{t-1,q}^{UK}} \right)$ is annual GVA growth in the UK relative to the same quarter in the previous year.⁴

The stacked VAR is a VAR (at the annual frequency) using $y_t = \left(y_{t,1}^{UK}, y_{t,2}^{UK}, y_{t,3}^{UK}, y_{t,4}^{UK}, y_t^A \right)'$ as the vector of dependent variables where $y_t^A = \left(y_t^{1,A}, \dots, y_t^{R,A} \right)$ stacks all the annual variables into vectors. In words, this approach stacks GVA growth for all the regions along with the four quarterly values for UK GVA growth into a vector which contains the dependent variables in a VAR.

The reduced form version of the stacked VAR⁵ with P lags is written as:

$$y_t = B_0 + \sum_{j=1}^P B_j y_{t-j} + \varepsilon_t \quad (1)$$

where B_0 is a vector of intercepts. We consider homoskedastic and heteroskedastic versions of the model. The former assumes ε_t to be i.i.d. $N(0, \Sigma)$.

⁴Note that we are not approximating the percentage growth rate using log differences. The use of log differences would entail slight changes in our entropic tilting formulae and, in particular, to the weights in (19) below.

⁵The stacked VAR is often written as a structural VAR which imposes a sequential ordering on the high frequency variables (see, e.g., McCracken, Owyang and Sekhposyan, 2016). To do impulse response analysis, such an ordering is required. But for unconditional forecasting, it is acceptable to use an unrestricted reduced form (see the discussion of section 2.3 of Ghysels, 2016). In this paper, we use the stacked VAR to produce unconditional forecasts which are then entropically tilted. Hence, we work with this reduced form VAR.

The heteroskedastic version of this model uses the specification of Cogley and Sargent (2005) which replaces the Σ of the homoskedastic model by Σ_t which is written as:

$$\Sigma_t = AD_tA' \quad (2)$$

where A is a lower triangular matrix with ones on the diagonal. D_t is a diagonal matrix with diagonal elements σ_{it}^2 which are assumed to follow univariate stochastic volatility processes. That is,

$$\sigma_{it}^2 = \exp(h_{it}) \quad (3)$$

where

$$h_{it} = \gamma h_{it-1} + v_{it} \quad (4)$$

with v_{it} i.i.d. $N(0, \phi_i)$. The homoskedastic specification is obtained if $\phi_i = 0$ for all i .⁶

Note that we are working with annual data which means the sample will be short. And, since the dimension of y_t is $N = R + 4$ we are working with a fairly large VAR. With large VARs such as this, it is common to use Bayesian methods so as to allow for prior shrinkage to overcome the problems associated with a shortage of data information.

2.2 Bayesian Analysis with the Stacked VAR

With large VARs, Bayesian methods using the Minnesota prior are commonly used (see, among many others, Banbura, Giannone and Reichlin, 2010) and we follow this practice with our mixed frequency VAR. However, following Giannone, Lenza and Primiceri (2015), we estimate the prior shrinkage parameters from the data. In this sub-section, we provide details (see also Dieppe, Legrand and van Roye, 2016, section 3.3).

We begin with the homoskedastic version of the model. The Minnesota prior replaces Σ by $\hat{\Sigma}$ which is the OLS estimate from the stacked VAR. Thus, we need only worry about the prior for the VAR coefficients. Let β be the $N \times (NP + 1)$ vector containing all the VAR coefficients. The Minnesota prior is $N(\underline{\beta}, \underline{V})$ with particular choices for $\underline{\beta}$ and \underline{V} . These can be explained by noting that the VAR coefficients can be divided into three categories: i) own lags (i.e. lags of dependent variable i in equation i), ii)

⁶The assumption of Normal errors is a standard one in the large Bayesian VAR literature when working with relatively short macroeconomic data sets such as ours. It is computationally efficient in that it allows for analytical posterior results under the Minnesota prior. Our inclusion of stochastic volatility allows for a greater degree of flexibility in modelling the dispersion of the errors in a time-varying manner.

other lags (i.e. lags of dependent variable i in equation j for $i \neq j$) and iii) exogenous variables such as the intercept. The prior mean vector, $\underline{\beta}$, is set to zero except for first own lag coefficients which are set to b . We consider a grid of values within the interval $b \in [0.1, 1.0]$ with a step size of 0.05 and estimate b .

The prior covariance matrix, \underline{V} , is a diagonal matrix with diagonal elements specified as follows:

- Prior variances for coefficients on own lags at lag l are:

$$\left(\frac{\lambda_1}{l^{\lambda_3}}\right)^2. \quad (5)$$

- Prior variances for coefficients on the l^{th} lag of the j^{th} variable in the i^{th} equation are:

$$\left(\frac{s_i^2}{s_j^2}\right) \left(\frac{\lambda_1 \lambda_2}{l^{\lambda_3}}\right)^2. \quad (6)$$

- The prior variance for the intercept is:

$$s_i^2 (\lambda_1 \lambda_4)^2. \quad (7)$$

In these expressions, s_i^2 is the OLS estimate of the error variance from a univariate autoregressive model for the i^{th} variable.

We estimate the shrinkage parameters. For λ_1 , which controls overall shrinkage, we use the grid of values in the interval $[0.05, 0.3]$ with a step size of 0.01. For λ_2 which controls other lag shrinkage, we use the grid of values in the interval $[0.1, 3]$ with a step size of 0.05. For λ_3 which controls the rate that shrinkage increases on longer lag lengths, we use a grid of values in the interval $[1, 2]$ with step size 0.2. For λ_4 we use a grid over the interval $[100, 1000]$ with step size 100 which implies a very non-informative prior. All these intervals contain the benchmark recommendations of Dieppe, Legrand and van Roye (2016) within them; and we did not obtain estimates at any of the boundaries of our grids indicating that they are sufficiently wide.

For the heteroskedastic version of the model we use the prior just described for the VAR coefficients, but additionally require a prior for the parameters controlling Σ_t . These are A , γ and ϕ_i and h_{i0} for $i = 1, \dots, N$. We set $\gamma = 0.85$ and let each free element of the lower triangular matrix A have a non-informative prior. For ϕ_i we use relatively non-informative inverse Gamma priors:

$$IG(0.001, 0.001). \quad (8)$$

As a general comment about prior specification, we have done extensive experimentation with various choices from the range of priors available in the BEAR Toolbox of Dieppe, Legrand and van Roye (2016). We have also experimented with different lag lengths. The specification and prior choices used in this paper are those which yield the highest marginal likelihoods. This led us to work with the Minnesota prior and set the lag length, P , to one.

Posterior and predictive analysis can be done using standard Bayesian MCMC methods and we use the BEAR toolbox to do so (see Dieppe, Legrand and van Roye, 2016). The main output will be draws from the one-step (and two-step) ahead predictive densities. For future reference, we will denote the predictive density of $y_{\tau+1}$ given all the information available at time τ by $p(y_{\tau+1}|Data_{\tau})$, where $Data_{\tau}$ denotes all the data available to the forecaster at the end of period τ . Given the aforementioned publication lags associated with the regional data, such that in the UK regional GVA data for year τ are not currently available until near the end of year $\tau + 1$, the predictive density of interest, $p(y_{\tau+1}|Data_{\tau})$, is in effect produced as a two-year ahead forecast from the stacked VAR until the regional data for year τ are published in December of year $(\tau + 1)$. That is, until late in year $(\tau + 1)$, rather than contain data for year τ , $Data_{\tau}$ in fact contains regional GVA data dated year $(\tau - 1)$ and earlier.

2.3 Entropic Tilting Using Quarterly Releases of UK Data

The previous sub-section described how to produce unconditional (with respect to current year information) forecasts using annual data. Given the (as of the time of writing this paper) nearly one year delay in releasing regional GVA data, these forecasts, $p(y_{\tau+1}|Data_{\tau})$, can be used as nowcasts for the year. However, we want to update these nowcasts throughout year $(\tau + 1)$ as new information on UK GVA is released each quarter. We will do so using entropic tilting methods as described in this sub-section.

The standard stacked VAR defined by (1) captures the general property that quarterly GVA growth data for the UK as a whole might help nowcast regional GVA growth, since lags of UK GVA growth appear on the right hand side of the equation for each region and the VAR error covariance matrix allows for contemporaneous correlations between the equations for regional GVA growth and that of the UK as a whole. This structure means that if we update UK GVA figures as they are released after each quarter, the regional GVA growth figures will also be updated. If, for instance, an unexpectedly

favourable outcome for UK GVA growth occurs in the first quarter of a year, this is a strong signal that growth in most or all UK regions has also increased. It is desirable to incorporate this information now (i.e. after the first quarter value of UK GVA has been released) and update the estimates of regional GVA throughout the year rather than waiting for the release of regional GVA data. The interlinkages built into the VAR allow us to do this.

But there is a second way that the quarterly releases of UK GVA data can be used to shed light on what is happening in the regions. This is through what we will call the cross-sectional restriction. This restriction embodies the fact that GVA growth for the UK as a whole is a weighted average of regional growth rates. In this sub-section, we discuss how to incorporate these types of information in the context of the stacked VAR using entropic tilting methods.

Increasingly, macroeconomic forecasters want to move beyond unconditional forecasts to incorporate extra information or restrictions on their forecasts (see, among many others, Alessi, Ghysels, Onorante, Peach and Potter, 2014, and Clark, Krueger and Ravazzolo, 2017) such as our cross-sectional restriction. Conditional forecasting and entropic tilting are two ways of doing this.⁷ The idea of conditional forecasting is that you impose this condition

⁷Anticipating the discussion in the remainder of this section, an alternative method for imposing restrictions goes back at least as far as Deming and Stephan (1941); with Bryon (1978) and Smith, Weale and Satchell (1998) developing the statistical theory. The underlying principles (e.g. typically involving least squares methods and focusing on the mean as opposed to the entire predictive density) differ from and are less general than entropic tilting, which also differs given its Bayesian foundations. Entropic tilting can also be used on other quantiles or moments of the distribution. Entropic tilting, in our context, delivers (restricted or posterior) regional estimates that minimise the (KLIC) distance between the unrestricted density $p(y_{\tau+1}|Data_{\tau})$ and the restricted density $p^*(y_{\tau+1}|Data_{\tau})$

such that the mean of $\sum_{r=1}^R w_{r,\tau} y_{\tau+1}^{r,A}$, using the restricted $y_{\tau+1}^{r,A}$ estimates, equals $y_{\tau+1,1}^{UK}$.

While the least squares methods deliver (restricted) regional estimates that minimise the quadratic distance between $y_{\tau+1,1}^{UK}$ and $\sum_{r=1}^R w_{r,\tau} y_{\tau+1}^{r,A}$ subject to the (exact) constraint that

$y_{\tau+1,1}^{UK} = \sum_{r=1}^R w_{r,\tau} y_{\tau+1}^{r,A}$. Assuming as we do that $p(y_{\tau+1}|Data_{\tau})$ and $p^*(y_{\tau+1}|Data_{\tau})$ are

Normal, the predictive mean of the entropically tilted (restricted) density, $p^*(y_{\tau+1}|Data_{\tau})$, is such that $E\left(\sum_{r=1}^R w_{r,\tau} y_{\tau+1}^{r,A}\right) = y_{\tau+1}^{r,A}$ whereas under least squares $\sum_{r=1}^R w_{r,\tau} y_{\tau+1}^{r,A} = y_{\tau+1}^{r,A}$.

We also note that entropic tilting with Normal predictive densities keeps the variance of the regional estimates unchanged, see (11); but least squares involves a reduction in the variance (cf. Bryon, 1978, equation 6).

exactly on the forecasts.⁸ This is increasingly done by policymakers in, for example, central banks. For instance, the policymaker may be interested in forecasts of inflation for different interest rate paths. An unrestricted VAR for inflation, the interest rate and other variables would provide unrestricted forecasts of inflation. Conditional forecasting procedures would allow, for instance, for a forecast of inflation conditional on the interest rate remaining at 0.5%, another forecast conditional on the interest rate being raised to 0.75%, etc. Conditional forecasting methods impose the restriction exactly. That is, in a predictive simulation algorithm which provides draws (call them $y_{\tau+1}^{(s)}$ for $s = 1, \dots, S$) from the predictive density, every single draw will satisfy the restriction. This contrasts with entropic tilting where only the predictive mean (or other predictive moments specified by the researcher) will satisfy the constraint. In this paper we use entropic tilting since we expect the cross-sectional restriction to hold only approximately and, thus, we do not wish to impose it exactly as in (“hard”) conditional forecasting (or least squares methods). In our case, this cross-sectional relationship is approximate since the GVA data for the regions that we use do not exactly add up to UK GVA because of measurement error (see the Data Appendix) and because our main results exclude GVA produced in the UK continental shelf (UKCS). UKCS data are dominated by the activities of the UK oil and gas sector.

As Table 5 shows, the UKCS data exhibit volatile behaviour that is also inconsistent with how the other regions relate to UK GVA. As a result, for our main results, we do not include UKCS in y_t^A for fear of contaminating the relationship between the other UK regions and UK GVA with potentially deleterious effects on the accuracy of the nowcasts. However, as it is ultimately an empirical matter what works best, we also present results which do include UKCS in y_t^A . In both cases, UKCS remains part of $y_{t,q}^{UK}$; i.e. the UK GVA figures that we condition the regional nowcasts on include the UKCS. This means that for those VARs that exclude UKCS in y_t^A this is an additional reason, to measurement error, why we expect the cross-sectional relationship to hold only approximately. Note that it is not possible to remove UKCS activity from the overall estimates of UK quarterly GVA and

⁸This type of conditional forecast is what Waggoner and Zha (1999) call a “hard” forecast. They also propose a “soft” methodology that, rather than imposing a value exactly, allows for some variability around this by centering a distribution on the value. As reviewed in Dieppe, Legrand and van Roye (2016), “soft” conditioning relates to entropic tilting. Entropic tilting has the advantage of both letting one impose conditions on any moment of the forecast density and, as discussed below, it imposes the condition in a way that ensures the conditional density forecasts are “as close as possible” to the original forecasts. Below we also show how entropic tilting methods can be extended, in our context, to impose cross-sectional (aggregation) constraints.

then entropically tilt towards that estimate. While some sectoral detail for GVA is available for the UK as a whole on a more timely basis, not all Oil and Gas related activity in the UK ‘Mining & quarrying including oil and gas extraction’ sector is activity which takes place in the UKCS. Some of this activity relates to onshore activity in support of activity in the UKCS. Similarly, not all of the activity in this sector relates to oil and gas extraction. It would therefore not be appropriate to treat the ‘Mining & quarrying including oil and gas extraction’ sector as synonymous with the UKCS activity series.

The idea of entropic tilting is to produce a new predictive density, $p^*(y_{\tau+1}|Data_\tau)$, which has a mean which satisfies the restriction but is in all other respects as close as possible to $p(y_{\tau+1}|Data_\tau)$. “As close as possible” is defined according to the Kullback-Leibler Information Criterion (KLIC) which is a measure of the relative entropy of $p^*(y_{\tau+1}|Data_\tau)$ to $p(y_{\tau+1}|Data_\tau)$. So, in our case, the predictive mean (i.e. the point forecast) produced by $p^*(y_{\tau+1}|Data_\tau)$ will satisfy the restrictions but otherwise the predictive density will be as close as possible to the unrestricted predictive density produced by the stacked VAR.

We use results based on a Normal approximation.⁹ Assume that the unrestricted predictive density is Normal:

$$y_{\tau+1}|Data_\tau \sim N(\mu, V) \quad (9)$$

and we break down the parameters into UK and regional blocks as follows:

$$\mu = \begin{bmatrix} \mu_{UK} \\ \mu_R \end{bmatrix}, V = \begin{bmatrix} V_{UK} & V'_{UK,R} \\ V_{UK,R} & V_R \end{bmatrix}. \quad (10)$$

The estimation procedure of the preceding sub-section will provide μ and V .

Now suppose that we want to tilt the predictive density so that the mean of some variables is fixed (e.g. so as to set the predictive mean of $y_{\tau+1}^{UK}$ to μ_{UK}^* where μ_{UK}^* is chosen to reflect period $\tau + 1$ UK-wide information that has come available before the $\tau + 1$ regional data are released), but otherwise we want to leave the predictive density to be as close to $p(y_{\tau+1}|Data_\tau)$ as possible. It can be shown (see, e.g., Altavilla, Giacomini and Ragusa, 2017) that the tilted predictive density is:

$$y_{\tau+1}^*|Data_\tau \sim N(\mu^*, V^*) \quad (11)$$

⁹Conditional on the parameters of the model, the predictive density from our model is Normal. The unconditional predictive density integrates out the parameters and, thus, is no longer Normal but is likely to be nearly so.

where $V^* = V$ (i.e. tilting does not change the predictive variance) and

$$\mu^* = \begin{bmatrix} \mu_{UK}^* \\ \mu_R - V_{UK,R}V_{UK}^{-1}(\mu_{UK} - \mu_{UK}^*) \end{bmatrix} = \begin{bmatrix} \mu_{UK}^* \\ \mu_R^* \end{bmatrix}. \quad (12)$$

Note that this type of entropic tilting relates to UK variables since this is what is being released throughout the year. Thus, it may appear that it does not directly impact on the regional growth nowcasts. But this appearance is incorrect since $\mu_R \neq \mu_R^*$. The intuition is that, unless $V_{UK,R} = 0$ and the UK nowcasts are uncorrelated with the regional nowcasts, the updating of UK GVA nowcasts will spill over into the regional nowcasts.

But we also want to tilt toward the cross-sectional constraint which does directly relate to the regional growth nowcasts. To add the latter restriction, we extend the conventional result given in (11). To this end, we define a new variable $z = Ay_{t+1}$. The properties of the multivariate Normal distribution imply

$$z \sim N(A\mu, AVA') \quad (13)$$

for any $M \times N$ matrix A . If we set

$$A = \begin{bmatrix} w \\ I_N \end{bmatrix} \quad (14)$$

where $w = (0, 0, 0, 0, w_{1,t-1}, \dots, w_{R,t-1})$ and $w_{r,t-1}$ for $r = 1, \dots, R$ are region-specific weights to be defined below, then z contains the weighted average of the nowcasts of regional GVA growth as its first element, followed by the four quarterly UK GVA growth nowcasts, followed by the R regional nowcasts.

We apply the entropic tilting formula of (11) to z . To this end, let $\mu^\dagger = A\mu$ and $V^\dagger = AVA'$ where

$$\mu^\dagger = \begin{bmatrix} \mu_1^\dagger \\ \mu_2^\dagger \end{bmatrix}, V^\dagger = \begin{bmatrix} V_{11}^\dagger & V_{21}^{\dagger'} \\ V_{21}^\dagger & V_{22}^\dagger \end{bmatrix} \quad (15)$$

and assume that the tilting restrictions are $\mu_1^\dagger = \mu_1^*$. Let z^\ddagger denote the tilted version of z . Then the same derivations used to find (11) can be used to show that:

$$z^\ddagger \sim N(\mu^\ddagger, V^\ddagger) \quad (16)$$

where

$$\mu^\ddagger = \begin{bmatrix} \mu_1^* \\ \mu_2^\dagger - V_{21}^\dagger V_{11}^{\dagger-1}(\mu_1^\dagger - \mu_1^*) \end{bmatrix}. \quad (17)$$

Note that V^\ddagger will be a singular matrix, but this causes no problem for our derivations as they only involve inverting V_{11}^\dagger (which is non-singular) and

we are only interested in the tilted predictive densities for the regional GVA variables which have predictive covariance matrix V_{22}^\dagger (which is non-singular).

The preceding material described the general motivation and formulae relating to entropic tilting. To describe the precise way we implement it (i.e. the exact choice for μ_1^*), we first define the temporal and cross-sectional constraints we will use. These results arise from the fact that annual UK GVA, Y_t^{UK} , can be written in two different ways:

$$Y_t^{UK} = \sum_{q=1}^4 Y_{t,q}^{UK} = \sum_{r=1}^R Y_t^{r,A}. \quad (18)$$

In growth rates, this implies

$$\begin{aligned} y_t^{UK} &= \frac{Y_t^{UK} - Y_{t-1}^{UK}}{Y_{t-1}^{UK}} = \sum_{q=1}^4 \left(\frac{Y_{t-1,q}^{UK}}{\sum_{q=1}^4 Y_{t-1,q}^{UK}} \right) y_{t,q}^{UK} \\ &= \sum_{r=1}^R \left(\frac{Y_{t-1,q}^{r,A}}{\sum_{r=1}^R Y_{t-1,q}^{r,A}} \right) y_t^{r,A} \\ &= \sum_{q=1}^4 w_{q,t-1}^{uk} y_{t,q}^{UK} = \sum_{r=1}^R w_{r,t-1} y_t^{r,A} \end{aligned} \quad (19)$$

where $y_{t,q}^{UK}$ is UK growth relative to the previous year and the w 's are the weights.

Now imagine we know $y_{t+1,1}^{UK}$ i.e. UK growth in the first quarter of year $(t+1)$. We wish to impose this information when nowcasting, but $y_{t+1,2}^{UK}, y_{t+1,3}^{UK}, y_{t+1,4}^{UK}$ are still unknown. We therefore assume, when tilting to reflect the cross-sectional constraint, that $y_{t+1,2}^{UK} = y_{t+1,3}^{UK} = y_{t+1,4}^{UK} = y_{t+1,1}^{UK}$ i.e. growth continues through year $t+1$ at the rate seen in the first quarter.¹⁰ This implies we tilt to reflect:

$$y_{t+1,1}^{UK} = \sum_{r=1}^R w_{r,t} y_{t+1}^{r,A}. \quad (20)$$

Now assume we know $y_{t+1,1}^{UK}$ and $y_{t+1,2}^{UK}$ and again assume growth continues at the most recent quarterly rate through the remainder of the year. This means we now tilt to reflect:

$$\left(w_{1,t}^{uk} y_{t+1,1}^{UK} + 3w_{2,t}^{uk} y_{t+1,2}^{UK} \right) = \sum_{r=1}^R w_{r,t} y_{t+1}^{r,A}. \quad (21)$$

¹⁰Given that our data are seasonally adjusted, the assumption of constant growth throughout the year is the most reasonable one and, as we shall see, it works well empirically.

Noting that the first element of μ_1^* will relate to the variable $\sum_{r=1}^R w_{r,t} y_{t+1}^{r,A}$, the following summarises how we proceed as we update our nowcasts using entropic tilting as new UK data (Q1 to Q4, i.e. $y_{\tau+1,1}^{UK}$ to $y_{\tau+1,4}^{UK}$) are released:

1. After the release of Q1 UK GVA growth (in May of each year) set

$$\mu_1^* = \left(y_{\tau+1,1}^{UK}, y_{\tau+1,1}^{UK} \right)'$$

2. After Q2 release (in August of each year) set

$$\mu_1^* = \left(\left(w_{1,\tau}^{uk} y_{\tau+1,1}^{UK} + 3w_{2,\tau}^{uk} y_{\tau+1,2}^{UK} \right), y_{\tau+1,1}^{UK}, y_{\tau+1,2}^{UK} \right)'$$

3. After Q3 release (in November of each year) set

$$\mu_1^* = \left(\left(w_{1,\tau}^{uk} y_{\tau+1,1}^{UK} + w_{2,\tau}^{uk} y_{\tau+1,2}^{UK} + 2w_{3,\tau}^{uk} y_{\tau+1,3}^{UK} \right), y_{\tau+1,1}^{UK}, y_{\tau+1,2}^{UK}, y_{\tau+1,3}^{UK} \right)'$$

4. After Q4 release (in February of each year) set

$$\mu_1^* = \left(y_{\tau+1,1}^{UK}, y_{\tau+1,1}^{UK}, y_{\tau+1,2}^{UK}, y_{\tau+1,3}^{UK}, y_{\tau+1,4}^{UK} \right)'$$

In fact, since prior to release of the Q4 data for year $\tau + 1$ the regional GVA data for year τ are not yet available (since 2005 the regional data have been published in December of each year), and so as to respect this publication lag and produce genuinely real-time nowcasts, we condition our Q1 to Q3 nowcasts on 2-year rather than 1-year ahead (unconditional) density forecasts from the VARs. While this not affect how we condition the regional nowcasts on within-year ($\tau + 1$) data for the UK, as detailed in (1) to (4) above, for the Q1 to Q3 nowcasts we consider an augmented A matrix and an augmented μ^\dagger vector, see (22) below, that let us impose the additional cross-sectional constraint that the regional data for year τ , while now forecast rather than assumed known as in Q4, are consistent with known UK data for (the previous) year τ that are available from when the Q1 nowcast is made for year $\tau + 1$.

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & w'_{\tau-1} & 0_{1 \times R} \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0_{1 \times R} & w'_\tau \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0_{1 \times R} & 0_{1 \times R} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & I_R & 0_R \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0_R & I_R \end{bmatrix} \begin{bmatrix} y_{\tau,1}^{UK} \\ y_{\tau,2}^{UK} \\ y_{\tau,3}^{UK} \\ y_{\tau,4}^{UK} \\ y_{\tau+1,1}^{UK} \\ y_{\tau+1,2}^{UK} \\ y_{\tau+1,3}^{UK} \\ y_{\tau+1,4}^{UK} \\ y_\tau^A \\ y_{\tau+1}^A \end{bmatrix} ; \mu^\dagger = \begin{bmatrix} y_\tau^{UK} \\ y_{\tau,1}^{UK} \\ y_{\tau,2}^{UK} \\ y_{\tau,3}^{UK} \\ y_{\tau,4}^{UK} \\ y_{\tau+1}^{UK} \\ y_{\tau+1,1}^{UK} \\ y_{\tau+1,2}^{UK} \\ y_{\tau+1,3}^{UK} \\ y_{\tau+1,4}^{UK} \\ y_\tau^A \\ y_{\tau+1}^A \end{bmatrix} \quad (22)$$

3 Empirical Results

In this section, we examine the performance of our nowcasting methods using data from 1967-2016. We use quarterly UK GVA growth data and annual GVA growth data for either 9 UK regions or 10 if we include UKCS. Definitions of these regions and further details of the data are given in the Data Appendix. It is worth reiterating that in this empirical exercise we use, as closely as possible, first release GVA estimates in our model, and compare our nowcasts to these same data. In this way our empirical exercise is as near as possible real-time. The key question of interest is whether the entropic tilting using timely, quarterly, UK-wide data will improve the nowcasts of (ONS first release) regional GVA data. This is the question we will focus on in this section.

To evaluate the nowcasts from the VAR models, we use a variety of standard measures of forecast performance. In particular, we use root mean squared forecast error (RMSFE) to evaluate the qualities of the point nowcasts. The evaluation of the accuracy of the entire predictive density uses log predictive scores (LPS) and the continuous ranked probability score (CRPS). See Appendix A.10 of Dieppe, Legrand, and van Roye (2016) for definitions of all these nowcast (forecast) evaluation metrics. We report these results for the VAR models relative to the (2-year ahead density) forecasts from AR(1) models (with Normal errors). These models simply take the annual regional data for each region individually and use ordinary least squares (OLS) methods; but given the aforementioned publication lags we use 2-year ahead forecasts. Comparison against a univariate benchmark enables us to

assess the utility in our VAR models of conditioning the regional nowcasts on within-year UK data exploiting inter-regional dynamics.

To aid in interpretation, note that relative values for the CRPS and RMSFE measures less than unity indicate that there are forecast gains associated with use of our VAR models; these *relative* values are calculated as the CRPS or RMSFE from our nowcasting model divided by the CRPS or RMSFE from the benchmark model. For the LPS we subtract the LPS from the benchmark model from those for each of our nowcasting models; positive values now therefore indicate improved forecast accuracy, relative to the univariate benchmark. Our nowcast evaluation period begins in 2006. Our methods are recursive. That is, we do a real-time out-of-sample nowcasting exercise using an expanding window of data beginning in 2006.

Tables 1 and 2 present these three forecast metrics, relative to the AR model, for the 9 UK regions for the homoskedastic and heteroskedastic stacked VARs. Tables 3 and 4 repeat the analysis using 10 regions, where UKCS is included in the stacked VAR. The final row of each panel of each table presents an average (for RMSFE and CRPS) or sum (for LPS) over all regions. As one moves from left to right in the tables, the forecasting metrics reflect more and more information. The first column of numbers in each table is based on unconditional nowcasts (i.e. 2-year density forecasts from the VARs). In all these tables, it can be seen that, except on one occasion if interested in nowcasting the UKCS, incorporating new information on UK GVA (as it accumulates each quarter) via our entropic tilting methods, works. That is, substantial decreases in RMSFEs and CRPSs and increases in LPSs, relative to the AR benchmark, are observed as we move through the year. For instance, in Table 1 we find that the average (across regions) of the RMSFEs is almost half as small by the end of the year as it was at the beginning (i.e. it drops from 0.97 to 0.49 as we move through the year). Thus, overall, the point forecasts are improving substantially. The LPS results show that similar improvements occur for the entire predictive density. For instance, in Table 1 the sum of the LPSs over all regions increases from 0.79 to 6.54.¹¹ The gains are also strong using the CRPSs with, in Table 1, the average over all regions dropping from 0.93 to 0.48 and the average CRPS dropping from 0.85 to 0.55 in Table 2. Interestingly most of the gains in nowcast performance are found after the first quarter of UK GVA data are released. Thus, nowcasts produced as early as April by our stacked VAR approach are

¹¹The ratio of predictive scores comparing any two models is similar to a Bayes factor (except that the Bayes factor is an evaluation of predictive performance over the entire sample). With Bayes factors a common rule of thumb (see Kass and Raftery, 1995) is that there is strong evidence in favour of one model over another if the the log Bayes factor is greater than 3. We are at or above this strong evidence threshold.

appreciably better than the unconditional nowcasts that would have been produced in January. There are nevertheless modest gains seen, on average in Table 1, as quarterly information accrues through the year with the RMSE and CRPS ratios declining and the LPS differences increasing. Interestingly, conditioning on the Q4 release does not help much; this is despite the fact that it is only with this Q4 release of UK GVA that the regional data for the previous year become available, so that we can condition on a 1 rather than a 2-year (unconditional) density forecast from the VAR.

The fact that our four tables are producing similar results offers reassurance that our results are robust to changes in specification and in data. We note that there is little evidence that inclusion of stochastic volatility is important in this application. It is true that, if we use conventional model comparison measures using the unconditional forecasts, the inclusion of stochastic volatility does lead to slight improvements relative to the homoskedastic model. For instance, the sum of the log predictive scores is 1.04 in Table 2 (which includes stochastic volatility) and 0.79 in Table 1 (which does not). A similar pattern can be found if we compare Tables 3 and 4. However, when we look at the entropically tilted nowcasts, the homoskedastic version of the model tends to do better. For instance, the sum of log predictive scores using tilted nowcasts with 4 quarters of UK GVA growth is 6.54 in Table 1 but only 4.54 in Table 2.

A comparison of Tables 3 and 4 with Tables 1 and 2 indicates that including UKCS as a region does not tend (with a few exceptions) to lead to any improvements in nowcast performance for the 9 other UK regions. For instance, the best overall summary of evidence is probably the sum of the log predictive scores for the 9 UK regions and for this a comparison of Tables 1 and Tables 3 indicates that including UKCS leads to a rise of 0.12 in the unconditional forecast case, but reductions in the log predictive scores across the other four nowcasts (relative to the AR benchmark). On this basis, we conclude that omitting UKCS is not harmful and take the homoskedastic mixed frequency VAR with nine regions as being our preferred specification to look at in more detail. We are not surprised by this result, given the distinct (univariate) time series properties of UKCS relative to the other regions of the UK as summarised in Table 5.

If we look at the individual regions, they uniformly exhibit the same patterns noted above. As new information about UK GVA is released (on a quarterly basis) it clearly is helping to improve nowcasts for every region. These (relative) nowcast improvements are particularly large for London and the South East. This is not surprising, since this region comprises a large share of UK GVA. But even for smaller regions (e.g. Scotland) we are finding nowcast improvements which are similarly large. Our weakest relative perfor-

mance is for Northern Ireland. But even for this region, the nowcast metrics improve as new information on UK GVA growth is incorporated throughout the year and are clearly more accurate than those from the AR benchmark (with gains of at least 20% across the three evaluation metrics).

It can be seen that while there are often gains from using an unconditional multivariate forecasting method which allows information from different regions and the UK as a whole to inform the forecasts of a particular region, our unconditional forecasting metrics (i.e. without entropic tilting) do not always beat the AR(1) benchmark. London and the South East is a leading example in Table 1. It is only when entropic tilting is used that we see the more substantive gains. That is, when we use the additional quarterly UK data as it released throughout the year large gains are made relative to a simple univariate method.

Tables 1 to 4 reflect average performance over our nowcast evaluation period. Figures 1 to 3 shed light on whether there are particular time periods when incorporating new information using our tilting methods is particularly important. These figures are based on the homoskedastic mixed frequency stacked VAR and do not include the UKCS. They plot, for each region, actual regional GVA growth (i.e. the subsequent realisation using the first estimate from the ONS) along with our five different nowcasts (conditional means of the nowcast densities). One clear pattern which emerges, to varying degrees across regions, is that 2009 is a year where updating regional nowcasts, in the light of the more timely UK data, is particularly important. As the Great Recession hit, the unconditional forecast of 2009 GVA growth turned out to be much higher than the realisation in every region. However, our tilting methodology quickly downgraded the 2009 nowcasts. 2013 is another year when the tilted nowcasts showed big improvements relative to the unconditional forecasts. In this case, realised GVA growth in all regions was higher than expected and the unconditional forecast was too low in every region. By tilting towards UK-wide releases as they came available, the nowcasts were upgraded and ended up being much closer to the actual 2013 realisation.

4 Conclusions

In this paper we have highlighted the need for more timely macroeconomic data for the UK regions. Our desire is to produce regional GVA estimates which are more frequent (quarterly instead of annual) and also more timely. We have developed an econometric procedure which combines a mixed frequency VAR with entropic tilting. Our key contributions lie in the incorporation of the new, more timely, information provided by the quarterly releases

of UK wide data acknowledging the fact that UK growth is a weighted average of growth for the individual regions. Exploiting this cross-sectional constraint, and noting that we do not expect it to hold exactly, we are able to produce updated regional nowcasts to the same timescale as the ONS currently produce their quarterly UK estimates. That is, the latest UK data help allocate national growth among the regions of the UK. In a real-time nowcasting exercise we find our methods to work well. As new, quarterly UK wide information is released throughout the year our nowcasts of regional GVA growth improve. Thus, using the methods we propose, regional policymakers can have at their disposal more accurate nowcasts of current growth rates. They do not have to rely on out-of-date figures or indeed have to wait many months, for new regional data releases from ONS.

Table 1: Nowcasting Performance Using Homoskedastic Mixed Frequency VAR (Results Relative to AR Benchmark). *NOTE: the RMSFE and CRPS values from our VAR nowcasting model are presented relative to (divided by) those from the benchmark AR model; in the same way the LPS are presented relative (by subtraction) to the LPS from the AR model.*

Tilting Using New Information:	None	Q1	Q2	Q3	Q4
RMSFE					
North	0.96	0.61	0.58	0.56	0.59
York. & Humber	0.92	0.48	0.51	0.49	0.48
East Mids	0.90	0.56	0.44	0.42	0.45
West Mids	0.88	0.47	0.37	0.35	0.37
Lon & SE	1.09	0.37	0.35	0.39	0.38
South West	0.97	0.48	0.45	0.42	0.35
Wales	1.06	0.62	0.55	0.53	0.58
Scotland	0.92	0.45	0.41	0.41	0.44
N. Ireland	1.05	0.71	0.77	0.77	0.76
Average RMSE	0.97	0.53	0.49	0.48	0.49
LPS					
North	0.23	0.57	0.60	0.61	0.71
York. & Humber	0.28	0.62	0.60	0.61	0.76
East Mids	0.08	0.54	0.63	0.65	0.76
West Mids	0.17	0.63	0.70	0.70	0.82
Lon & SE	-0.21	0.71	0.73	0.70	0.80
South West	0.16	0.59	0.61	0.63	0.76
Wales	0.08	0.57	0.63	0.64	0.73
Scotland	0.14	0.65	0.68	0.68	0.86
N. Ireland	-0.14	0.43	0.35	0.35	0.35
Sum LPS	0.79	5.32	5.53	5.58	6.54
CRPS					
North	0.89	0.58	0.56	0.55	0.55
York. & Humber	0.85	0.52	0.53	0.52	0.47
East Mids	0.86	0.53	0.45	0.44	0.43
West Mids	0.80	0.48	0.41	0.41	0.39
Lon & SE	1.15	0.44	0.43	0.45	0.42
South West	0.90	0.51	0.49	0.47	0.40
Wales	1.01	0.60	0.54	0.53	0.55
Scotland	0.88	0.47	0.44	0.44	0.43
N. Ireland	1.06	0.69	0.74	0.73	0.69
Average CRPS	0.93	0.53	0.51	0.50	0.48

Table 2: Nowcasting Performance Using Mixed Frequency VAR with Stochastic Volatility (Results Relative to AR Benchmark) *NOTE: the RMSFE and CRPS values from our VAR nowcasting model are presented relative to (divided by) those from the benchmark AR model; in the same way the LPS are presented relative (by subtraction) to the LPS from the AR model.*

Tilting Using New Information:	None	Q1	Q2	Q3	Q4
RMSFE					
North	0.78	0.63	0.61	0.59	0.63
York. & Humber	0.71	0.45	0.40	0.38	0.37
East Mids	0.75	0.55	0.44	0.43	0.43
West Mids	0.82	0.51	0.40	0.39	0.39
Lon & SE	0.92	0.36	0.34	0.39	0.36
South West	0.79	0.47	0.44	0.41	0.37
Wales	0.86	0.58	0.52	0.50	0.56
Scotland	0.76	0.46	0.41	0.41	0.43
N. Ireland	0.88	0.69	0.75	0.74	0.75
Average RMSE	0.81	0.52	0.48	0.47	0.48
LPS					
North	0.06	0.16	0.22	0.23	0.40
York. & Humber	0.20	0.32	0.36	0.37	0.60
East Mids	0.17	0.33	0.40	0.41	0.59
West Mids	0.02	0.31	0.38	0.38	0.51
Lon & SE	-0.04	0.33	0.34	0.33	0.49
South West	0.12	0.32	0.34	0.35	0.50
Wales	0.18	0.35	0.39	0.39	0.51
Scotland	0.23	0.42	0.42	0.42	0.59
N. Ireland	0.09	0.22	0.20	0.20	0.35
Sum LPS	1.04	2.76	3.04	3.09	4.54
CRPS					
North	0.87	0.74	0.73	0.73	0.67
York. & Humber	0.81	0.66	0.63	0.62	0.51
East Mids	0.74	0.61	0.54	0.53	0.46
West Mids	0.87	0.64	0.57	0.57	0.51
Lon & SE	1.00	0.62	0.61	0.63	0.54
South West	0.83	0.63	0.61	0.60	0.51
Wales	0.86	0.68	0.63	0.63	0.59
Scotland	0.77	0.57	0.55	0.55	0.50
N. Ireland	0.89	0.73	0.77	0.76	0.70
Average CRPS	0.85	0.65	0.63	0.62	0.55

Table 3: Nowcasting Performance Using Homoskedastic Mixed Frequency VAR including UKCS (Results Relative to AR Benchmark) *NOTE: the RMSFE and CRPS values from our VAR nowcasting model are presented relative to (divided by) those from the benchmark AR model; in the same way the LPS are presented relative (by subtraction) to the LPS from the AR model.*

Tilting Using New Information:	None	Q1	Q2	Q3	Q4
RMSFE					
North	0.96	0.54	0.52	0.52	0.67
York. & Humber	0.91	0.49	0.39	0.40	0.55
East Mids	0.92	0.68	0.54	0.50	0.42
West Mids	0.89	0.51	0.34	0.33	0.37
Lon & SE	1.09	0.55	0.42	0.43	0.45
South West	0.97	0.56	0.46	0.43	0.42
Wales	1.03	0.67	0.56	0.54	0.52
Scotland	0.91	0.46	0.37	0.39	0.46
N. Ireland	1.05	0.76	0.71	0.72	0.76
UKCS	3.69	1.37	0.99	0.95	1.16
Average RMSE (Inc. UKCS)	1.24	0.66	0.53	0.52	0.58
Average RMSE (Exc. UKCS)	0.97	0.58	0.48	0.47	0.51
LPS					
North	0.24	0.62	0.63	0.63	0.59
York. & Humber	0.29	0.60	0.64	0.64	0.69
East Mids	0.06	0.40	0.54	0.57	0.79
West Mids	0.16	0.59	0.71	0.72	0.82
Lon & SE	-0.19	0.56	0.66	0.66	0.73
South West	0.16	0.53	0.59	0.61	0.72
Wales	0.13	0.50	0.59	0.60	0.82
Scotland	0.16	0.62	0.68	0.67	0.83
N. Ireland	-0.12	0.36	0.43	0.41	0.35
UKCS	0.50	0.65	0.67	0.67	0.77
Sum LPS (Inc. UKCS)	1.41	5.44	6.14	6.17	7.10
Sum LPS (Exc. UKCS)	0.91	4.78	5.47	5.50	6.34
CRPS					
North	0.89	0.53	0.52	0.53	0.61
York. & Humber	0.84	0.52	0.47	0.48	0.52
East Mids	0.88	0.63	0.52	0.49	0.40
West Mids	0.81	0.51	0.40	0.40	0.39
Lon & SE	1.14	0.58	0.48	0.49	0.48
South West	0.90	0.56	0.49	0.47	0.44
Wales	0.97	0.65	0.55	0.54	0.49
Scotland	0.87	0.48	0.43	0.44	0.44
N. Ireland	1.06	0.74	0.68	0.68	0.69
UKCS	1.06	0.73	0.70	0.70	0.65
Average RMSE (Inc. UKCS)	0.94	0.59	0.52	0.52	0.51
Average RMSE (Exc. UKCS)	0.93	0.58	0.50	0.50	0.50

Table 4: Nowcasting Performance Using Mixed Frequency VAR with Stochastic Volatility and including UKCS (Results Relative to AR Benchmark) *NOTE: the RMSFE and CRPS values from our VAR nowcasting model are presented relative to (divided by) those from the benchmark AR model; in the same way the LPS are presented relative (by subtraction) to the LPS from the AR model.*

Tilting Using New Information:	None	Q1	Q2	Q3	Q4
RMSFE					
North	0.79	0.59	0.59	0.58	0.65
York. & Humber	0.72	0.42	0.36	0.35	0.42
East Mids	0.76	0.60	0.50	0.49	0.44
West Mids	0.83	0.52	0.42	0.41	0.41
Lon & SE	0.92	0.42	0.34	0.35	0.36
South West	0.79	0.46	0.42	0.37	0.37
Wales	0.85	0.62	0.56	0.54	0.55
Scotland	0.76	0.46	0.37	0.38	0.48
N. Ireland	0.89	0.72	0.71	0.71	0.77
UKCS	1.82	0.53	0.90	0.98	0.90
Average RMSE (Inc. UKCS)	0.91	0.53	0.52	0.51	0.54
Average RMSE (Exc. UKCS)	0.81	0.53	0.47	0.46	0.49
LPS					
North	0.11	0.19	0.22	0.23	0.41
York. & Humber	0.23	0.31	0.35	0.35	0.61
East Mids	0.19	0.27	0.33	0.34	0.62
West Mids	0.09	0.27	0.34	0.33	0.52
Lon & SE	0.04	0.28	0.30	0.31	0.52
South West	0.15	0.29	0.31	0.33	0.54
Wales	0.18	0.30	0.34	0.35	0.56
Scotland	0.22	0.39	0.42	0.42	0.57
N. Ireland	0.08	0.20	0.21	0.21	0.35
UKCS	0.44	0.46	0.46	0.46	0.63
Sum LPS (Inc. UKCS)	1.74	2.96	3.27	3.33	5.33
Sum LPS (Exc. UKCS)	1.30	2.50	2.81	2.88	4.70
CRPS					
North	0.86	0.72	0.72	0.72	0.68
York. & Humber	0.79	0.65	0.62	0.62	0.52
East Mids	0.75	0.65	0.59	0.58	0.46
West Mids	0.86	0.67	0.60	0.59	0.51
Lon & SE	0.99	0.66	0.64	0.63	0.52
South West	0.82	0.63	0.61	0.59	0.49
Wales	0.86	0.71	0.67	0.66	0.57
Scotland	0.78	0.57	0.54	0.54	0.52
N. Ireland	0.90	0.76	0.74	0.73	0.70
UKCS	0.91	0.85	0.86	0.87	0.73
Average RMSE (Inc. UKCS)	0.85	0.69	0.66	0.65	0.57
Average RMSE (Exc. UKCS)	0.85	0.67	0.64	0.63	0.55

Table 5: Descriptive statistics for annual regional nominal GVA growth rates (1967-2016)

	Mean	Standard Deviation	Correlation with UK GVA
North	0.0774	0.0531	0.9409
York. & Humber	0.0790	0.0569	0.9157
East Midlands	0.0844	0.0620	0.9236
London & South East	0.0875	0.0485	0.9464
South West	0.0889	0.0621	0.9086
West Midlands	0.0777	0.0499	0.9374
Wales	0.0791	0.0590	0.9230
Scotland	0.0819	0.0549	0.9522
N. Ireland	0.0892	0.0643	0.7915
UKCS	-0.2589	2.8419	-0.1863

Figure 1: Plots of GVA Growth, $y_t^{r,A}$, and Nowcasts for the UK Regions

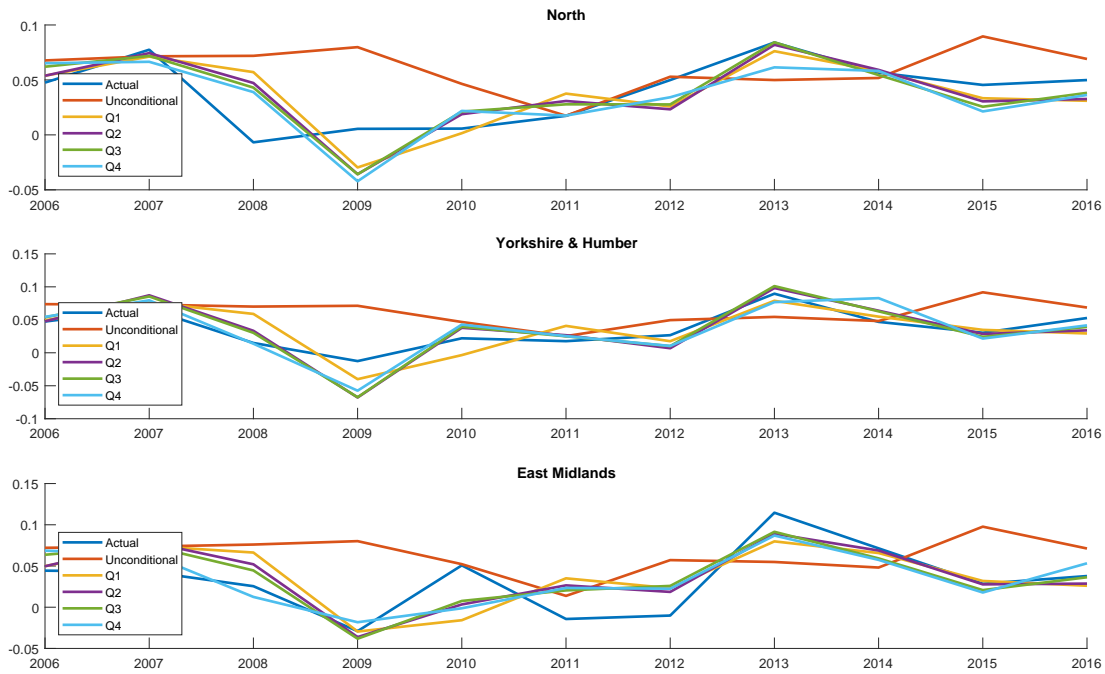


Figure 2: Plots of GVA Growth, $y_t^{r,A}$, and Nowcasts for the UK Regions

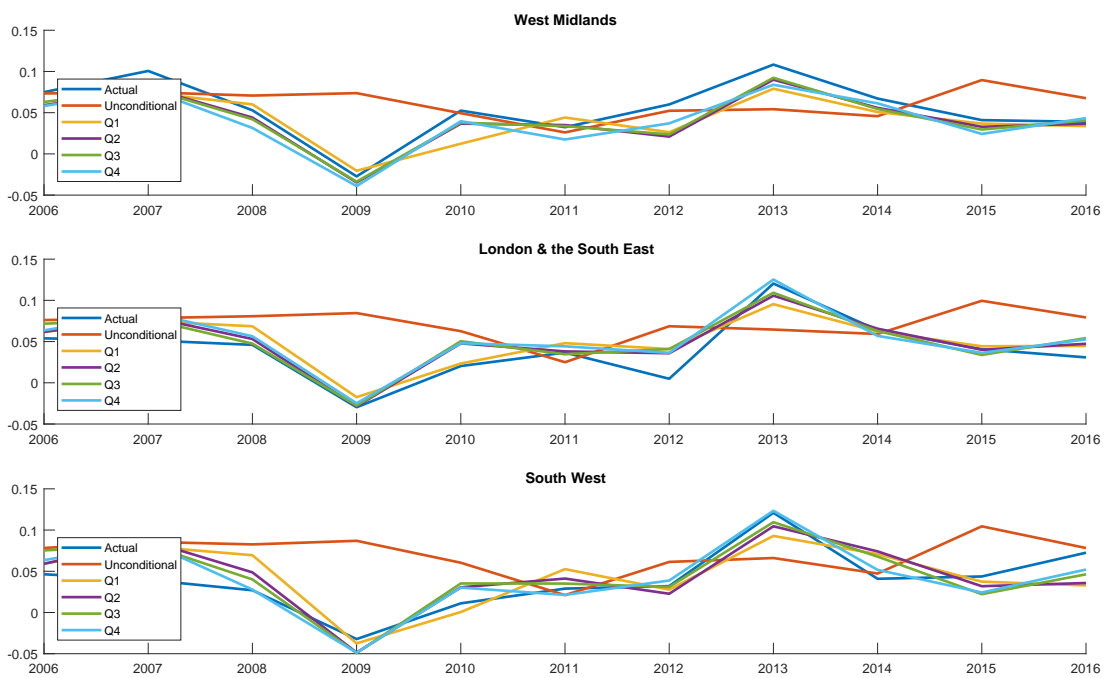
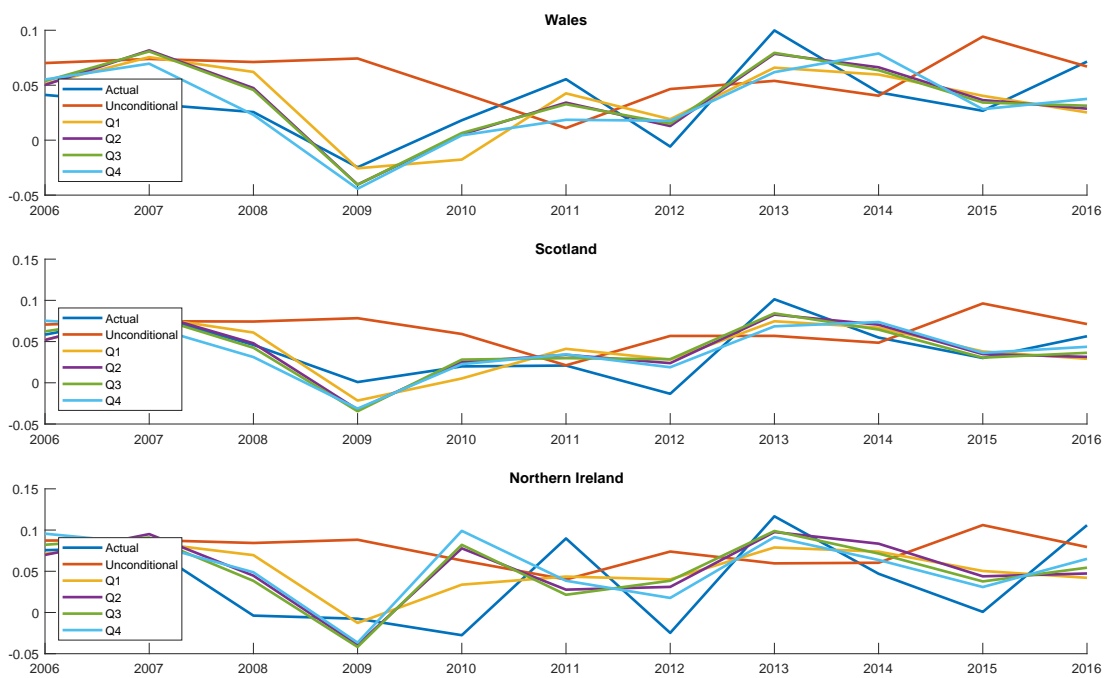


Figure 3: Plots of GVA Growth, $y_t^{r,A}$, and Nowcasts for the UK Regions



References

- Alessi, L., Ghysels, E., Onorante, L., Peach, R. and Potter, S. (2014). “Central bank macroeconomic forecasting during the global financial crisis: The European Central Bank and Federal Reserve Bank of New York experiences,” *Journal of Business and Economic Statistics*, 32, 483-500.
- Altavilla, C., Giacomini, R. and Ragusa, G. (2017). “Anchoring the yield curve using survey expectations,” *Journal of Applied Econometrics*, forthcoming.
- Banbura, M., Giannone, D. and Reichlin, L. (2010). “Large Bayesian vector autoregressions,” *Journal of Applied Econometrics*, 25, 71-92.
- Brave, S., Butters, R. and Justiniano, A. (2016). “Forecasting economic activity with mixed frequency Bayesian VARs,” *Federal Reserve Bank of Chicago Working Paper 2016-05*.
- Byron, R.P. (1978). “The estimation of large social account matrices,” *Journal of the Royal Statistical Society: Series A*, 141, 359-367.
- Carriero, A., Clark, T. and Marcellino, M. (2016). “Common drifting volatility in large Bayesian VARs,” *Journal of Business and Economic Statistics*, 34, 375-390.
- Clark, T. (2011). “Real-time density forecasts from BVARs with stochastic volatility,” *Journal of Business and Economic Statistics*, 29, 327-341.
- Clark, T., Krueger, F. and Ravazzolo, F. (2017). “Using entropic tilting to combine BVAR forecasts with external nowcasts,” *Journal of Business and Economic Statistics*, 35, 470-485.
- Clements, M. P., and Galvao, A. B. (2013). “Real-time forecasting of ináation and output growth with autoregressive models in the presence of data revisions,” *Journal of Applied Econometrics*, 28(3), 458-477
- D’Agostino, A., Gambetti, L. and Giannone, D. (2013). “Macroeconomic forecasting and structural change,” *Journal of Applied Econometrics*, 28, 82-101.
- Deming, W. and Stephan, F. (1941). “On a least squares adjustment of a sampled frequency table when the expected marginal totals are known,” *Annals of Mathematical Statistics*, 11, 427-444.
- Dieppe, A., Legrand, R. and van Roye, B. (2016). “The BEAR toolbox,” *European Central Bank working paper 1934*.
- Eraker, B., Chiu, C., Foerster, A., Kim, T. and Seoane, H. (2015). “Bayesian mixed frequency VAR’s,” *Journal of Financial Econometrics*, 13, 698-721.
- Ghysels, E. (2016). “Macroeconomics and the reality of mixed frequency data,” *Journal of Econometrics*, 193, 294-314.
- Ghysels, E., Grigoris, F. and Ozkan, N. (2017). “Forecasting of state and local government budgets: Exploiting mixed frequency and cross-border

data,” manuscript.

Giannone, D., Lenza, M. and Primiceri, G. E. (2015). “Prior selection for vector autoregressions,” *Review of Economics and Statistics*, 27, 436-451.

Kass, R. and Raftery, A. (1995). “Bayes Factors,” *Journal of the American Statistical Association*, 90, 773-795.

McCracken, M., Owyang, M. and Sekhposyan, T. (2016). “Real-time forecasting with a large, mixed frequency Bayesian VAR,” manuscript available at <http://www.tateviksekhposyan.org/>.

Mandalinci, Z. (2015). “Effects of monetary policy shocks on UK regional activity: A constrained MFVAR approach,” School of Economics and Finance, Queen Mary University of London, working paper 758.

Mikosch, H. and Neuwirth, S. (2015). “Real-time forecasting with a MIDAS VAR,” Bank of Finland Institute for Economies in Transition Discussion Paper 13-2015.

Primiceri, G. (2005). “Time varying structural vector autoregressions and monetary policy,” *Review of Economic Studies*, 72, 821-852.

Robertson, J., Tallman, E. and Whiteman, C. (2005). “Forecasting using relative entropy,” *Journal of Money, Credit and Banking*, 37, 383-401.

Schorfheide, F. and Song, D. (2015). “Real-time forecasting with a mixed-frequency VAR,” *Journal of Business and Economic Statistics*, 33, 366-380.

Smith, R., Weale, M. and Satchell, S. (1998). “Measurement error with accounting constraints: Point and interval estimation for latent data with an application to U.K. Gross Domestic Product,” *Review of Economic Studies*, 65, 109-134.

Waggoner, D. F. and Zha, T. (1999). “Conditional forecasts in dynamic multivariate models,” *Review of Economics and Statistics*, 81, 639-651.

Regional Data Appendix

This appendix summaries the data sources and construction of the regional GVA (income approach) database for the UK used in this paper. It describes the process of arriving at an annual dataset for nominal GVA for 9 ‘regions’ of the UK (plus the UK Continental Shelf) from 1966 to 2016 that is as consistent as possible, given changes to accounting standards over the time period. These changes mean that our regional estimates are measured at factor cost prior to 1996 and at basic prices from 1997.

A1 Data sources and matching against UK data

Our ambition in putting together the database was to use, as near as possible (certainly over our out-of-sample window), first-release estimates of regional GVA, at basic prices, and match these with the appropriate, similarly dated, data release for UK GVA. This strategy is in part motivated by our interest in nowcasting first release regional GVA estimates. But it also reflects the reality that final vintage data, e.g. the ONS’s latest regional estimates, are not available over the whole sample period (i.e. the latest ONS data, published in December 2017, cover the period 1997-2016 only). So to get earlier data we inevitably have to look to earlier data vintages. In matching the regional data to the UK data we sought to minimise the cross-sectional aggregation error, as ideally the sum of the regional GVA data equals the annual sum of the quarterly UK data. But, we should emphasise (as is detailed below) that it was not possible to eradicate this measurement error for all years. This motivates our use of tilting methods to approximately impose the cross-sectional aggregation constraint reflecting this measurement error.

The regional GVA data all come from the ONS (CSO) but via three sources:

1. The historical regional GDP database, recently published by the ONS, provides estimates, at factor cost, from 1966-1996, compiled from historical editions of the ‘Regional Trends’ and ‘Economic Trends’ journals.¹² The ONS Blue Book definition of factor cost states that “in the System of National Accounts 1968 this was the basis of valuation which excluded the effects of taxes on expenditure and subsidies”. By

¹²<https://www.ons.gov.uk/economy/regionalaccounts/grossdisposablehouseholdincome/adhocs/006226historiceconomicdataforregionsoftheuk1966to1996>

contrast, the latest ONS regional GVA estimates, considered in 2. and 3. below, are published in basic prices which exclude taxes (less subsidies) on products but do include taxes on the production process (such as business rates and any vehicle excise duty paid by businesses).

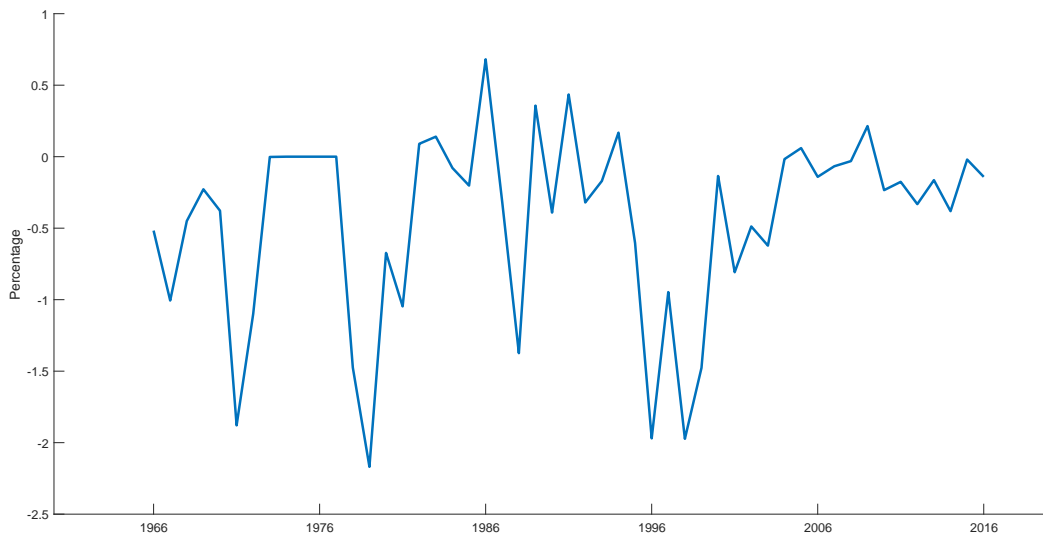
- The historical regional database “can be used as a proxy for the current regional GVA estimates”, as explained by the ONS in their supporting documentation. They also note that these data were “produced under the various statistical standards, regional and industry breakdowns which were current at the time they were first published”. The historical database does not always pick up estimates from successive yearly publications of Regional Trends. Our understanding, following email communication with ONS, is that this is because ONS chose to publish, in this historical database, the latest iteration for a given year rather than the first. As our interest is in extracting a database of first estimates, we deviate from the historical database as follows. From the (first) 1973 Regional Trends publication we extract regional data from 1966 to 1971. Thereafter, we consult successive annual Regional Trends publications so that the 1972 regional data come from the 1974 publication, the 1973 data come from the 1975 publication, and so on.
2. Successive annual issues of Economic Trends/Regional Trends (published in 1998 to 2005) were consulted to obtain regional GVA estimates, at basic prices, from 1997 to 2004.
 - This means the regional data are first release data.
 3. The GVA NUTS1 regional GVA revisions dataset is consulted to provide first release regional GVA estimates, at basic prices, from 2005-2016.¹³ These regional estimates are published with an eleven month lag, so that the 2005 data come from the December 2006 publication, and so on.

From 1966 to 1996 these regional data are matched against quarterly UK GVA data (at factor cost, seasonally adjusted) extracted from successive, similarly dated, national account data releases (obtained from the Bank of

¹³<https://www.ons.gov.uk/economy/grossvalueaddedgva/datasets/revisionstrianglesregionalgrossvalueaddedincomeapproachincurrentbasicprices>

England’s real-time database for nominal income; code CGCB¹⁴) with the secondary aim of minimising the cross-sectional aggregation measurement error of the sum of the regional data against the quarterly UK data when aggregated to the annual frequency. From 1997 the regional data are matched against successive, similarly dated (so that again the data vintages of the regional data match that of the UK data), releases of quarterly UK GVA estimates, at basic prices, from the ONS’s “Second estimate of GDP” previously known as the “UK Output, Income and Expenditure” press release/bulletins. Figure 1 shows that since 1997 (and our use of first release data) the cross-sectional aggregation measurement error is time-varying and not zero. The average statistical discrepancy between 1966 and 1996 is -0.47%, between 1997 and 2016 it is -0.39%¹⁵.

Figure A1: Discrepancy, by year, between the UK Quarterly series and Regional Annual series, as % UK GVA in each year



¹⁴Available at http://www.bankofengland.co.uk/statistics/Documents/gdpdatabase/nominal_income.xlsx

¹⁵It is worth noting that in the historical data that were released (data source 1 above) there was an explicit entry for ‘Statistical discrepancy’ and this accounts for the gap in Figure A.1 before 1997. In the later data no similar statistical discrepancy is formally reported, although as explained here, while small, a statistical discrepancy does emerge from a comparison of the first release of regional GVA data and the similarly dated (vintage) value of UK GVA data.

A2 Geographic reconciliation

The next step was to reconcile the different geographic breakdowns implied by the three data sources above. While the original Regional Trends publications (the historical data, 1. above) use “Standard Statistical Regions”, the later data sources for 1997 onwards (2. and 3. above) use NUTS1 regions.

This means that the historical data provide estimates for the following geographies: United Kingdom; North; Yorkshire and Humberside; East Midlands; East Anglia; South East; Greater London (1978 onwards); Rest of South East (1978 onwards); South West; West Midlands (2); North West; England; Wales; Scotland; Northern Ireland; United Kingdom Continental Shelf (UKCS). Prior to 1978 London was not separately identified in the regional data, instead it was part of the South East Standard Statistical Region. Between 1978 and the introduction of the NUTS classification system the old South East region was split into Greater London and Rest of South East. With the introduction of the NUTS classification, the Rest of South East region was split and one part merged with the old East Anglia Standard Statistical Region (which existed in the data from 1966-1994) to create a new ‘East of England’ NUTS1 region, and the other part maintained as the NUTS1 region ‘South East England’.

NUTS1 data are therefore presented for the following areas: United Kingdom; North East; North West & Merseyside; Yorkshire and the Humber; East Midlands; West Midlands; Eastern; London; South East; South West; England; Wales; Scotland; Northern Ireland; United Kingdom less Continental Shelf.

To arrive at a consistent series - for what we call our final dataset - we aggregated both geographical classifications to produce 9 (10 including the UKCS) “regional” GVA series which, we believe, are consistent in terms of geographic coverage across the two regional definitions. Table A1 details how the GVA series were aggregated across standard statistical and NUTS1 regions to arrive at our final “regions”.

Table A1: Regional definitions

Historical data	NUTS1	KMM Region ID	Final dataset name
North	North East	1	North
	North West & Merseyside	1	North
North West		1	North
Yorkshire and Humberside	Yorkshire and the Humber	2	Yorkshire and Humber
East Midlands	East Midlands	3	East Midlands
West Midlands	West Midlands	4	West Midlands
South East		5	London & South East
	Greater London (>1978)	5	London & South East
Rest of South East (>1978)	London	5	London & South East
	South East (GOR)	5	London & South East
East Anglia		5	London & South East
	Eastern	5	London & South East
South West	South West	6	South West
Wales	Wales	7	Wales
Scotland	Scotland	8	Scotland
Northern Ireland	Northern Ireland	9	Northern Ireland
United Kingdom CS	United Kingdom CS	10	United Kingdom CS

Looking at Table A1, we see the London and South East England was the most problematic region. This reflects the fact that at the beginning of the sample data are reported only for the South East of England (encompassing London and the rest of the South East, although not reported separately) and East Anglia (which was about 8% the size of the South East region in GVA terms). The difficulty is that we cannot disaggregate the South East, in the early part of these data, into London and the rest of the South East. In addition, were we able to do so, in practice the values for the South East Standard Statistical Region (pre-1995) do not align well with those for the NUTS1 region ‘South East’ in 1995.

Figure A2 illustrates the correspondence between the statistical regions and the Government Office (or NUTS1) regions. This map illustrates this difficulty that we encountered in the South East of England. The East Anglia region, as reported from 1966-1994 (in the historical data, 1. above) is not coterminous with the subsequent East of England NUTS1 region. Similarly, the old Standard Statistical Region ‘South East of England’ (1966-1994, although split out into ‘Greater London’ and ‘Rest of South East’ from 1978) includes parts of what is now ‘East of England’, as well as ‘London’ and the ‘South East of England’.

In addition, we can see that from Figure A2 that in the North of England, the ‘North’ Standard Statistical Region comprised parts of what, under

Figure A2: Government Office Region boundaries and Standard Statistical Region boundaries (<http://www.celsius.lshtm.ac.uk/modules/geog/ge030301.html>)



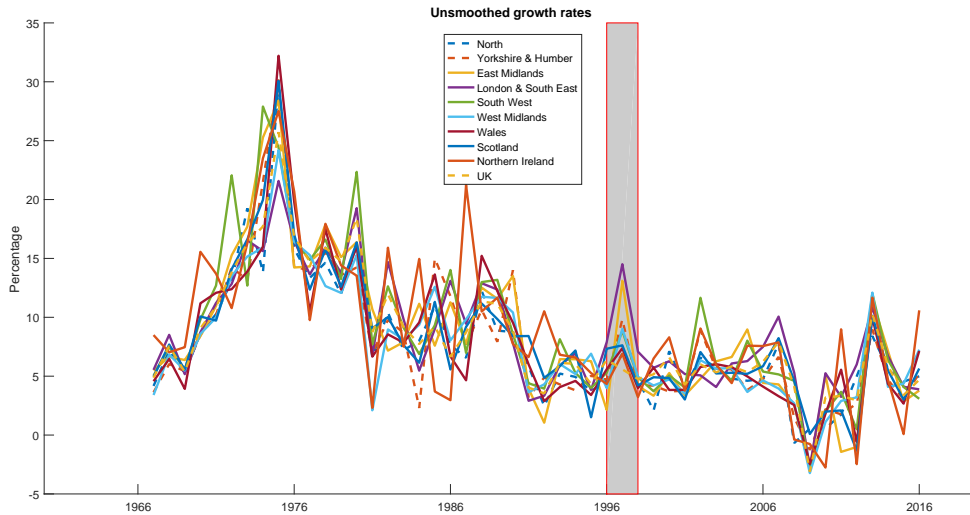
the NUTS1 classification, is now ‘North East’ and ‘North West’ regions of England.

Our strategy to derive a consistent database was therefore to aggregate both geographies to the most disaggregated common boundary. This results in the 9 “regions”, plus the UKCS, that we work with in the main paper.

A3 Data adjustments

Figure A3 plots annual nominal GVA growth rates (in %), for each of our 9 regions and for the UK as a whole, using these data. Inspection of this figure reveals a spike in growth between 1996 and 1997. To give an example, there is an 11% increase in UK GVA between 1996 and 1997, whereas separate ONS figures for the UK lead us to expect growth half this rate. Since we know this spike is, therefore, a feature of how we have merged the data from

Figure A3: Original regional nominal GVA growth rates (in %: $100 \times y_t^{r,A}$)



1. (above) with that from 2. and 3. (above), and reflects in particular the move from estimates at factor cost to basic prices, we treat it as an *outlier*; recall that 1996 to 1997 was also when the aforementioned change in how regions were measured took place, with an apparent (upward) level shift in the series. We therefore elected to smooth out this spike in the 1996-1997 annual growth rate. As our regional econometric models are estimated in annual growth rates, rather than (log) levels, our practical solution is simply to proxy the 1997 growth rate with the average of the growth rates in 1996 and 1998. This, in fact, brings the UK growth rate for 1996-1997 into line with that for ONS figures for UK GVA as a whole. Figure A4 presents the annual growth rates (in %) for each region having made this adjustment. We note that this adjustment falls outside the out-of-sample window we use to assess the nowcasting performance of our models.

Figure A4: Smoothed regional nominal GVA growth rates (in %: $100 \times y_t^{r,A}$)

