Bank of England

Nowcasting GDP at the Bank of England: a Staggered-Combination MIDAS approach

Macro Technical Paper No. 2
June 2025

Andre Moreira

A series designed to document models, analysis and conceptual frameworks for monetary policy preparation – they are written by Bank staff to encourage feedback and foster continued model development.

Macro Technical Paper Series

Dr Bernanke's 2024 **review** of the monetary policymaking processes at the Bank of England provided a number of constructive recommendations for reform which we are taking forward.

This included improving our model maintenance and development. Macroeconomic models and frameworks play an important role in the monetary policy process. To maximise the value of macroeconomic models, they must be well documented and continuously improved as time goes on.

This Macro Technical Paper (MTP) series is part of the Bank's response to Dr Bernanke's recommendations. These MTPs are authored by Bank staff, and are intended to document models, analysis, and conceptual frameworks that underpin monetary policy preparation. The models documented in the series will typically be used to assess the current state of the economy, forecast its future, and to simulate alternative paths and policy responses.

Importantly, while each MTP will provide insights about a particular model or modelling framework that is an 'input' to policy, no single model can possibly capture all the relevant features to perform even just one of those roles adequately. Models will inevitably have to be updated and will improve over time, including as they are adapted to different constellations of macroeconomic conditions. This is a natural part of real-time monetary policy making. So, inevitably, no MTP will provide definitive answers.

The Bank seeks to encourage an active and informed debate about its modelling frameworks. Publishing and discussing the analytical work undertaken to support its monetary policy choices is central to this ambition. These MTPs will support a culture of continuous learning in monetary policy making. As time goes on, Bank staff will update and upgrade models, drawing on insights from the frontier of the academic literature. Moreover, this transparency will encourage external engagement from experts to ensure our modelling tools remain fit for purpose.

Clare Lombardelli
Deputy Governor for Monetary Policy

Bank of England

Macro Technical Paper No. 2

Nowcasting GDP at the Bank of England: a Staggered-Combination MIDAS approach

Andre Moreira[(1)]

Abstract

This paper introduces a Staggered-Combination MIDAS (SC-MIDAS) approach, used by Bank of England staff to nowcast UK GDP and other variables. SC-MIDAS uses a mix of restricted and unrestricted MIDAS regressions and two forecast combination steps to exploit 'hard' and 'soft' data optimally through the release cycle – specifically when the lower frequency target (**quarterly** GDP) is also sampled at a higher frequency (**monthly** GDP). This structure enables it to capture key features of the data, including the mechanical relationship between monthly and quarterly GDP, and to dynamically reweight hard versus soft signals in a way that improves performance compared to standard pooled MIDAS approaches. In practice, SC-MIDAS combines accuracy with interpretability, outperforming several benchmarks out-of-sample and producing a range of outputs that policymakers find useful. Our positive experience with other UK variables shows SC-MIDAS is applicable more widely.

Key words: Nowcasting, MIDAS, forecast combination, GDP, Bank of England.

JEL classification: C53, E37.

---

(1) Bank of England. Email: andre.moreira@bankofengland.co.uk

I am particularly grateful to Fausto Grinspun for excellent analytical contributions. I would also like to thank Giulia Mantoan, Harvey Daniell and Tamarah Shakir for helpful comments and suggestions, and Current Economic Conditions colleagues for countless interactions on the development and use of the model. The views expressed in this paper are those of the author and do not necessarily represent those of the Bank of England or any of its committees.

The Bank's macro technical paper series can be found at
www.bankofengland.co.uk/macro-technical-paper/macro-technical-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 2978-3194 (online)

1 Introduction

Nowcasting is the science (sometimes, art) of 'predicting the present, the very near future, and the very recent past' (Bańbura et al., 2013). Nowcasts are useful because they exploit high-frequency information to provide a timelier read on the economy compared to more lagged official statistics. At the Bank of England, staff's assessment of current economic conditions is a key input to the Monetary Policy Committee's (MPC) decisions (Broadbent, 2023). Nowcasts of headline macroeconomic variables also feature regularly in external MPC communications such as the Monetary Policy Report (MPR) (Bank of England, 2025).

Over recent decades, nowcast models have become common currency at central banks like the ECB (Baffigi et al., 2004; Angelini et al., 2011) and the Fed (Higgins, 2014; Bok et al., 2018), with a variety of methods emerging in the literature (Ghysels and Marcellino, 2018). Several are also used at the Bank of England (Bell et al., 2014; Anesti et al., 2017). It is important to remember that published MPC projections are ultimately judgemental however, reflecting all available information, including from off-model sources. As such, there is no single model, or combination of models, that can be used to pre-empt them.

This paper introduces a Staggered-Combination MIDAS approach (SC-MIDAS), used by Bank of England staff to nowcast UK GDP and other variables. In addition to documenting a staple of the Bank's toolkit, we contribute to the literature by showing how a mix of restricted and unrestricted MIDAS regressions and two forecast combination steps can be used to exploit 'hard' and 'soft' data optimally through the release cycle – specifically when the lower-frequency target (*quarterly* GDP) is also sampled at a higher frequency (*monthly* GDP). The hard/soft data distinction has been explored extensively in dynamic factor models, where those are often treated as 'blocks' (Bańbura and Rünstler, 2011), but less so in the MIDAS literature, which has typically focussed on indicator-level pooling (Clements and Galvão, 2009). SC-MIDAS bridges that gap, with an approach tailored explicitly to hard and soft data as distinct informational groups that helps to improve performance.

SC-MIDAS was first developed in response to the introduction of 'Monthly GDP' by the Office for National Statistics (ONS, 2018; Bank of England, 2018). Prior to that, the Bank had relied on a three-model system, with different models better suited to different types of inputs and, by implication, different stages of the release cycle (Anesti et al., 2017). This created challenges, requiring staff to revisit their judgemental model weightings after each data release. The arrival of monthly GDP further underscored the need for a model that could handle both hard and soft data well through the data cycle – including the crucial mechanical link between monthly and quarterly GDP that some of our earlier models had missed. In doing so, SC-MIDAS has effectively automated many of those difficult judgements.

Figure 1: SC-MIDAS nowcasts at MPR forecast cut-off date, 2005-19

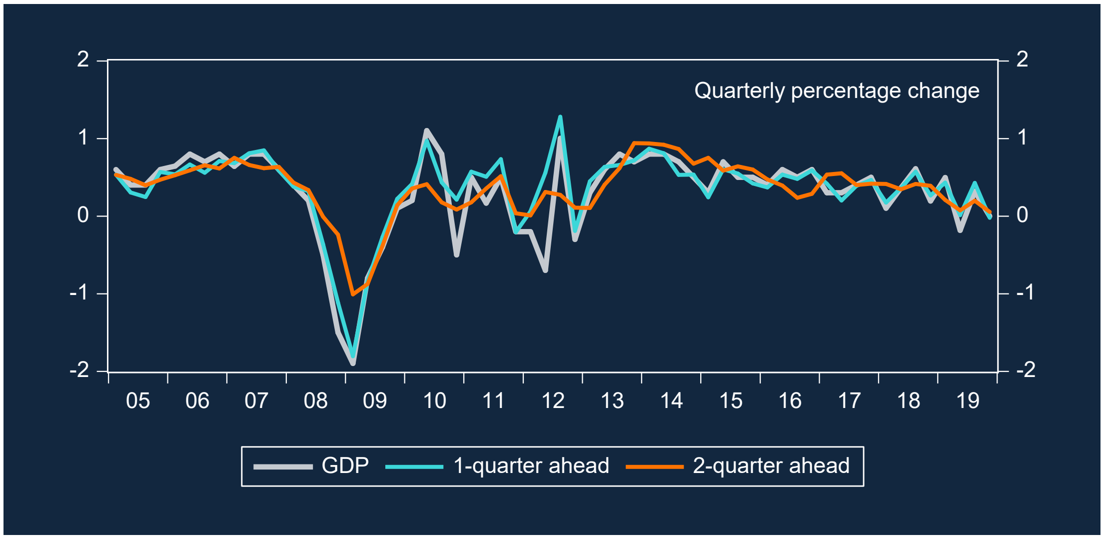

Notes: This figure shows one- and two-quarter-ahead SC-MIDAS nowcasts at the MPR cut-off date, around 25 days before quarterly GDP. Results are based on out-of-sample testing over 2005-19, using real-time-vintage monthly GDP data on the basis of the ONS' current publication schedule. The model performs well at both horizons compared to the series of historical GDP First Estimates.

In practice, SC-MIDAS combines accuracy with interpretability, outperforming several benchmarks in out-of-sample testing and producing a range of outputs that policymakers find useful. Figure 1 gives a preview of model performance at the MPR cut-off date, around 25 days before the quarterly GDP publication. One-quarter-ahead nowcasts are highly accurate at that stage, leveraging available monthly GDP data. Two-quarter ahead predictions are necessarily less precise but still anticipate broad trends and turning points well, informed by timelier surveys. Figures 2 and 3 show other examples of model outputs that the MPC see regularly, including several layers of SC-MIDAS nowcasts and their evolution through a typical 180-day nowcast window.

Methodological overview

At its core, nowcasting is about synthesising multiple, unsynchronised data releases into a single best estimate of the variable of interest (Giannone et al., 2008). SC-MIDAS addresses this in two stages: an initial 'MIDAS regression' stage that extracts independent signals from multiple indicators, and a 'staggered combination' procedure that aggregates them into the full SC-MIDAS nowcast (Figure 5). MIDAS is a 'direct forecasting' approach, relating available high-frequency lags to the lower-frequency target directly at the frequency and horizon of interest, so the exact same steps apply in the production of its one- and two-quarter ahead predictions.

Figure 2: SC-MIDAS nowcast layers, 2024Q4

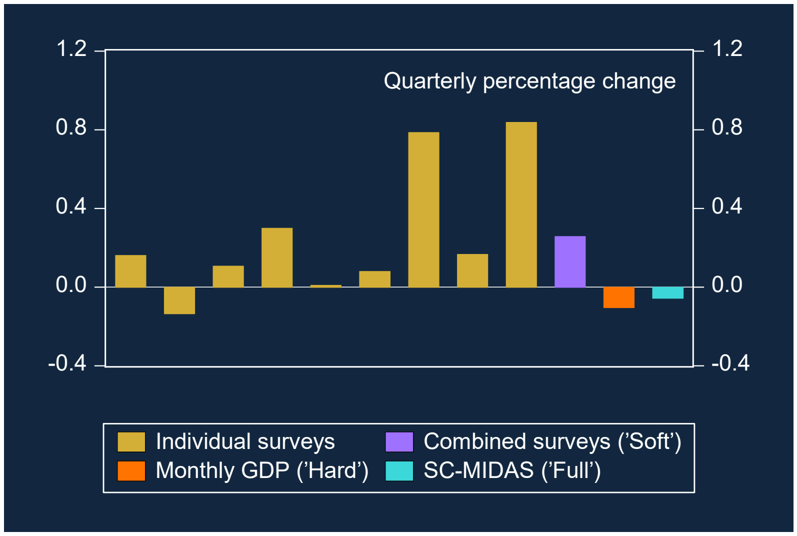

Notes: This figure shows layers of nowcasts produced as part of SC-MIDAS, including clean reads from each of the surveys, a combined 'soft' measure, a 'hard' measure based on monthly GDP, and the 'full' SC-MIDAS nowcast. The full range of nowcasts is often useful for policymakers to consider.

Unlike standard pooled MIDAS approaches, which treat all indicators symmetrically, SC-MIDAS is structured explicitly around the hard/soft data distinction. The example in this paper targets first estimates of quarterly GDP growth using 'hard' monthly GDP data and nine 'soft' survey indicators (Table 1). Each of these data types has different strengths and limitations. Monthly GDP is inherently informative about quarterly GDP growth given their 'time-aggregation' link – the fact that monthly and quarterly rates of the same concept must mechanically add up – but is also noisy and published with a 40-day lag. Surveys, by contrast, are timelier and smoother, but can also be highly collinear, requiring care when multiple indicators are involved.

The initial MIDAS stage involves separate regressions (Clements and Galvão, 2009) for each of monthly GDP and the surveys. MIDAS models are 'tightly parameterised, reduced form regressions that involve processes sampled at different frequencies' (Ghysels et al., 2004), where in practice a number of 'restricted' and 'unrestricted' specifications can be used (Ghysels et al., 2007; Foroni et al., 2015). For monthly GDP we strictly prefer unrestricted MIDAS, where the extra flexibility helps to better capture time aggregation mechanics (Figures 6, 7). For surveys the choice turns out to be less material, but we typically prefer a restricted Almon specification (Almon, 1965), offering a good balance between parsimony and computational ease.

The staggered combination stage aggregates the resulting set of predictions into the full SC-MIDAS nowcast in two further steps. While appropriate forecast combinations should generally reduce forecast error variance (Bates and Granger, 1969), more sophisticated methods can often

Figure 3: SC-MIDAS evolution through the release cycle, 2022Q4

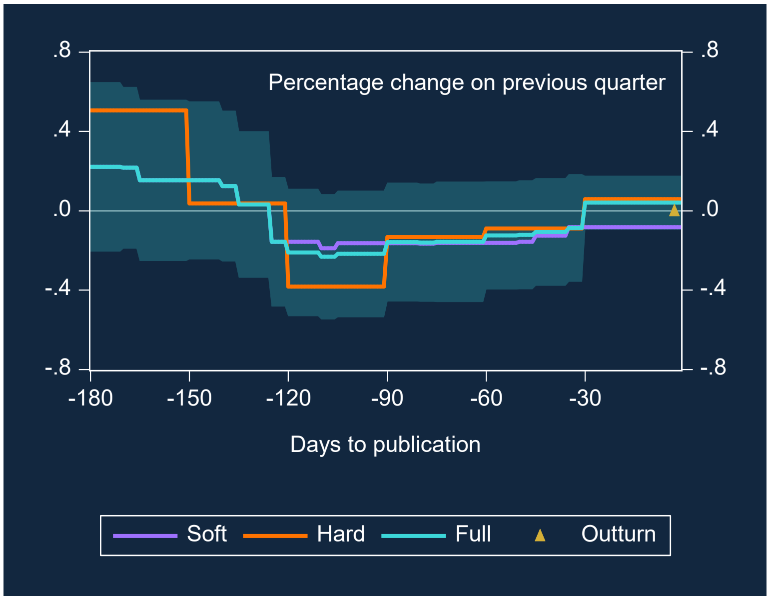

Notes: This figure shows the evolution of the 'full' SC-MIDAS nowcast and component 'hard' and 'soft' signals through the 180-day nowcast window. The full model closely tracks the soft nowcast through the first half of the window but increasingly matches up to the hard monthly GDP toward the end. Shaded band represents +/- 1 RMSE, based on 2005-19 out-of-sample testing.

fail in practice due to finite-sample estimation error (Smith and Wallis, 2009). Because surveys tend to be highly collinear, we first use a more robust 'RMSE-based' weighting to aggregate their nowcasts into a single combined 'soft' measure (Stock and Watson, 2004). The overall 'soft' and 'hard' signals are then combined into the 'full' SC-MIDAS nowcast in a final regression step, where in this case there is enough independent information to estimate their weights optimally by OLS (Granger and Ramanathan, 1984).

A further advantage of the SC-MIDAS setup is that it can easily accommodate full 'real-time-vintage' data, which we use for both estimation and evaluation (Koenig et al., 2003). This provides the fairest possible representation of model performance under real-world conditions. As is typical with MIDAS, the entire model is re-estimated for each new stage of the release cycle, allowing all regression parameters and combination weights to adapt.

Key findings

SC-MIDAS formalises many of Bank staff's existing practices, such as drawing on a wide range of indicators to inform the near-term outlook, or shifting weight flexibly between soft and hard signals to reflect their relative usefulness at different horizons (Anesti et al., 2017; Daniell and

Moreira, 2023). In our chosen application, SC-MIDAS places full weight on survey signals at the start of the nowcast window, but then gradually shifts up to 90% of that onto monthly GDP, with the publication of 'month 2' largely pinning the nowcast at the 30-day mark (Figure 9).

Nowcast accuracy improves markedly through the quarter as SC-MIDAS exploits signals from incoming data. Based on real-time out-of-sample testing over 2005-19, root mean squared errors (RMSE) drop from around 0.4 percentage points (pp) at 180 days, to 0.3pp at 90 days, and 0.1pp at 30 days (Figure 11). Within the detail, the full SC-MIDAS nowcast always performs at least as well as each of its hard and soft components, echoing the classic result from the forecast combination literature (Bates and Granger, 1969).

It also compares favourably in a 'horse-race' against several benchmarks – including a bridge model (Bridge, Bell et al., 2014), a dynamic factor model (DFM, Bok et al., 2018), and a more conventional pooled-MIDAS model (P-MIDAS, Anesti et al., 2017), similar to those in the Bank's old system. We show that SC-MIDAS is both capable of matching P-MIDAS and DFM's ability to exploit surveys at the start of the nowcast window, *and* of mimicking Bridge's mechanical time-aggregation of monthly GDP toward the end (Figure 12) – all within a single model. SC-MIDAS more generally outperforms, reducing RMSEs by up to 66%, 42% and 11% compared to DFM, P-MIDAS and Bridge (Table 2). Gains relative to the standard pooled MIDAS are particularly concentrated at the end of the quarter, when monthly GDP becomes the key determinant.

Our positive experience with other UK variables shows SC-MIDAS is applicable more widely (Daniell and Moreira, 2023). A preview of results for employment and CPI shows remarkably similar performance patterns (Figure 13). While SC-MIDAS was designed specifically to be used with a highly informative indicator like monthly GDP, it can in principle also handle more disaggregated hard data, by likewise pre-combining their signals before the final OLS step. Some of its advantage over standard pooled MIDAS might then erode though. A more general limitation relates to the number and variety of inputs that can be included, where the RMSE-based combination step in particular presupposes careful variable selection. For broader datasets, more complex shrinkage techniques may be preferred (Kohns and Potjagailo, 2023).

Practical use

SC-MIDAS is one of several models in the Bank's nowcasting suite. It plays a central role in informing GDP nowcasts and increasingly other variables'. Staff's GDP nowcasts nevertheless continue to draw on a plurality of models, including SC-MIDAS variants, monthly bridge equations (Bell et al., 2014) and dynamic factor models (Anesti et al., 2022). Manual adjustments are often made for one-off events, domestic and global developments, or intelligence from the Bank's 'Agents'.

Judgements about the informativeness of certain periods or variables can also be applied through model-based levers such as estimation windows, lag lengths, or input selections. The Covid period has been excluded permanently from estimation, for example. Post-pandemic model performance appears broadly intact (Figure 14).

Day to day, staff update SC-MIDAS for each new release, briefing the MPC as appropriate. The full model prediction is shown regularly, sometimes along with relevant 'nowcast layers' and weights for an intuitive decomposition. The degree of individual nowcast dispersion within SC-MIDAS can also be of interest as a gauge of near-term uncertainties. Related work by Bank colleagues has adapted a similar approach to a quantile setting, modelling different parts of the distribution directly (Mantoan and Verlander, 2025). In public, the MPC has sometimes found it helpful to cite SC-MIDAS' combined *soft* measure (the first step of staggered combination; Figure 6) as an 'indicator-based' proxy for the underlying trajectory of the economy in the face of more volatile official data (e.g. Broadbent, 2023; Mann, 2025; Ramsden, 2025).

The Bank's toolkit is under constant development. Staff will continue to provide updates on our preferred modelling approaches as part of a wider programme of improvements following the Bernanke review of forecasting for monetary policy making and communication at the Bank of England (Bernanke, 2024; Bank of England, 2024; Lombardelli, 2024).

Paper roadmap

The rest of the paper is structured as follows. Section 2 describes the UK GDP release cycle and data used in our application. Section 3 details the SC-MIDAS methodology and key estimation results. Section 4 presents findings from out-of-sample evaluation and horse-race exercises, along with brief extensions to other variables and the post-Covid period. Section 5 concludes.

2 Data

This section describes the UK GDP release cycle and inputs to the model. SC-MIDAS is set up to target 'first estimates' of quarterly GDP using monthly GDP and nine survey indicators. This relatively narrow dataset contains the bulk of exploitable signal for UK GDP in our experience, broadly matching that in staff's preferred version of the model. Because GDP is subject to revision, we follow best practice in using full 'real-time-vintage data' (Koenig et al., 2003). The distinction between 'hard' official data and 'soft' survey data (Bańbura and Rünstler, 2011) lies at the heart of the SC-MIDAS design.

Figure 4: Stylised calendar

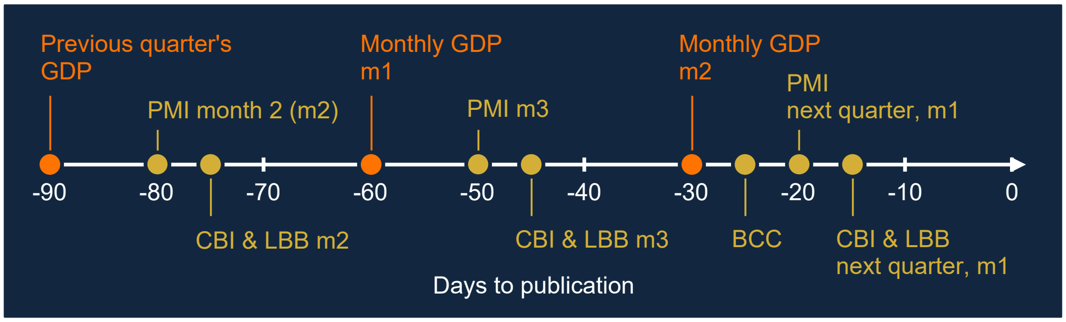

Notes: This diagram shows the stylised release calendar assumed for several results throughout the paper. This ensures results are representative of conditions faced by forecasters today.

2.1 Release cycle

We generally refer to the sequence of releases in the 180 days leading up to the quarterly GDP publication as the 'release cycle'. Results throughout the paper are based on a stylised quarterly release calendar, representative of the information flow that forecasters face *today*, rather than past release schedules. This ensures informational consistency across both estimation and evaluation, providing a more accurate picture of likely future performance.

Figure 4 shows the 'stylised calendar' with timings expressed in days to the quarterly GDP publication. For example, months 1 and 2 ('m1' and 'm2') of monthly GDP are assumed to be published at exactly 60 and 30 days to quarterly GDP, although real-world publication dates can vary slightly ('m3' always coincides with the quarterly publication). Monthly surveys are considerably timelier, with the equivalent 'PMIs' for example published at around 110 (m1) and 80 (m2) days. The quarterly 'BCC' survey is published just over a month before GDP.

2.2 Quarterly GDP

Quarterly GDP has been measured in the UK since the 1950s and remains a key economic statistic. At the Bank, the focus on calendar quarters remains even as the Office for National Statistics (ONS) has begun to publish monthly GDP data. This is partly a function of policymaker preferences, but also of the nature of our medium-term forecast machinery and processes, which continue to operate in quarterly space (Burgess et al., 2013; Albuquerque et al., 2025).

The ONS publishes its 'first quarterly estimates' of GDP growth with a lag of around 40 days relative to the end of the quarter. SC-MIDAS is specifically set up to target these 'first estimates', consistent with the Bank's tradition of treating data revisions separately (Cunningham et al., 2012) and more recent evidence suggesting that revisions to headline GDP have in any case become harder to predict (Robinson, 2016).

2.3 Monthly GDP

Compared to quarterly GDP, monthly GDP provides a more frequent but noisy read on UK activity (Appendix A). Monthly GDP is inherently informative about quarterly growth given the 'time-aggregation' link between the two – the fact that the monthly and quarterly rates must mechanically add up – but is published with the same long, 40-day lag.

For consistency with the left-hand-side variable, it is important to also use real-time vintages of monthly GDP on the right (Appendix B). While monthly GDP has only been published since 2018 (ONS, 2018), sectoral output information on which that is based has been available longer. To enable full real-time estimation of the model, we have extended existing monthly GDP vintages back to 2001, by aggregating the corresponding services, production and construction output series according to their contemporaneous GVA weights. This mirrors exactly how the ONS builds up monthly GDP today. Because monthly construction data only began in 2014, a slightly narrower proxy is used prior to that.

2.4 Surveys

Surveys are known to be among the best leading indicators of UK GDP growth, with a track record of signalling turning points (Mitchell, 2009). Unlike other countries where GDP is estimated from the expenditure side, the ONS bases its first estimates exclusively on output information (ONS, 2016). This creates a clearer conceptual link to business surveys, whose precision is nevertheless limited by the fact that these are qualitative 'net balances' – simple summary measures of directional ('up', 'down', 'same') responses. Surveys are fundamentally useful because they are both timelier and smoother than official data, but are also highly collinear (Appendix A), requiring care in estimation when multiple indicators are involved.

The nine indicators used in this application are sourced from widely monitored business surveys from S&P Global (PMI), the Confederation of British Industry (CBI), British Chambers of Commerce (BCC), and Lloyds Bank (LBB). The PMI, CBI and LBB surveys are particularly timely, published before the end of the reference month. The quarterly BCC is only published around the end of the calendar quarter. All included metrics relate to firm output, current or expected. These generally move together because they are conceptually similar metrics. Regardless of their precise stated horizon, measures of output *expectations* are also useful for nowcasting

Table 1: Survey indicator details

Category	Indicator	Provider	Survey	Frequency
Current output	PMI output	S&P Global	Composite PMI	Monthly
	CBI output past 3m	CBI	CGI & sectoral*	Monthly
	BCC sales	BCC	QES**	Quarterly
Near-term exp.	PMI new orders	S&P Global	Composite PMI	Monthly
	CBI output next 3m	CBI	CGI & sectoral	Monthly
	BCC new orders	BCC	QES	Quarterly
12m expectations	PMI future output	S&P Global	Composite PMI	Monthly
	BCC turnover exp.	BCC	QES	Quarterly
	LBB activity exp.	Lloyd's Bank	Business Barometer	Monthly

Notes: This table provides detail on included surveys, grouped by question horizon. *CGI = Composite Growth Indicator; extended to 1998 using sectoral CBI surveys (Service Sector, Industrial Trends, Distributive Trades). **QES = Quarterly Economic Survey; QES measures are calculated from services & manufacturing components.

and particularly so around turning points. Table 1 provides further detail on included surveys, grouped by question (not forecast) horizon.

3 Model

This section provides further detail on the method and key estimation results. SC-MIDAS is designed to exploit 'hard' and 'soft' data optimally through the release cycle, specifically in settings where the lower-frequency target (*quarterly* GDP) is also sampled at higher frequency (*monthly* GDP). It proceeds in two stages: an initial 'MIDAS regression' stage extracting independent signals from multiple indicators, and a 'staggered combination' stage aggregating those into the full SC-MIDAS nowcast. Figure 5 provides a visual summary. Unlike standard pooled-MIDAS approaches, which treat all indicators symmetrically, SC-MIDAS is structured deliberately around the hard/soft data distinction and the different accuracy/overfitting trade-offs that they pose. As is typical of MIDAS (Ghysels et al., 2004), the entire model is re-estimated at each stage of the release cycle, allowing all regression parameters and combination weights to adapt. This adaptability is key to its robust performance through the entire release cycle. Estimation throughout uses full back-runs of data, with Covid quarters (2020Q1-2021Q4) excluded through period dummies.

Figure 5: SC-MIDAS schematic

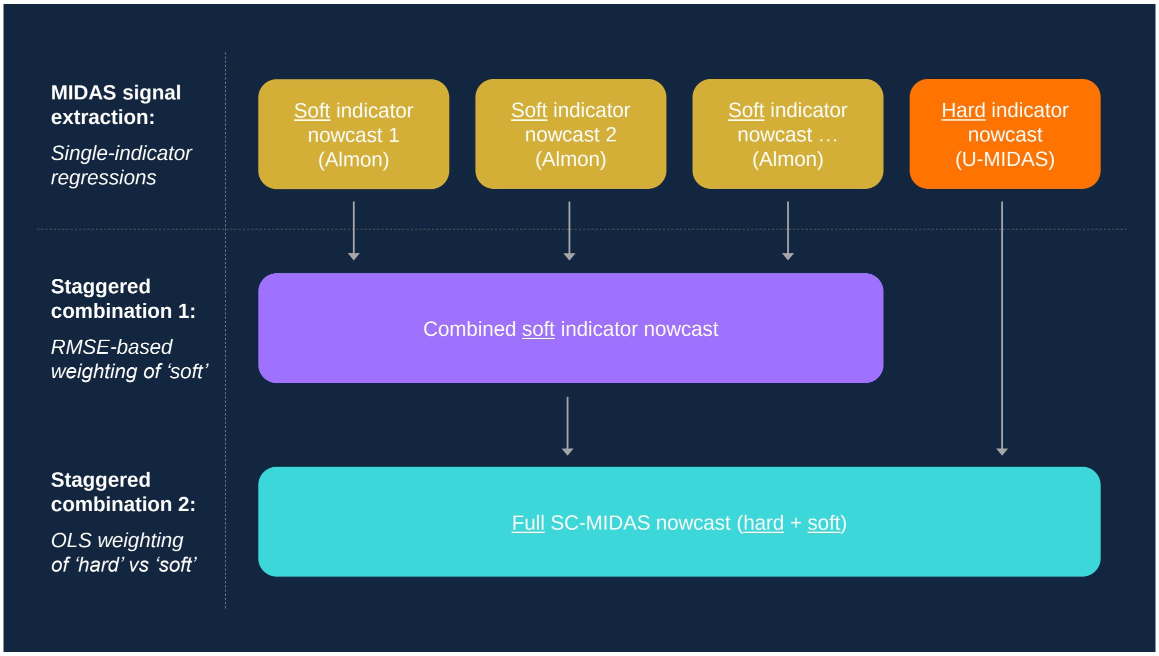

Notes: This diagram summarises the two stages of SC-MIDAS: a 'MIDAS regression' stage extracting independent signals from 'soft' and 'hard' inputs, and a 'staggered combination' aggregating those into the 'full' nowcast.

3.1 MIDAS regressions

The initial MIDAS regression stage extracts signals from each of the monthly GDP and survey indicators (Clements and Galvão, 2009) using a mix of 'unrestricted' and 'restricted' specifications. This variable-by-variable approach automatically deals with the 'ragged edge' (Giannone et al., 2008) of the dataset and sidesteps collinearity issues that would otherwise emerge from including several surveys in the same regression. It also produces a range of independent nowcasts that aid model interpretability.

MIDAS models are 'tightly parameterised, reduced form regressions that involve processes sampled at different frequencies' (Ghysels et al., 2004). The standard representation is:

$$y_t^q = \beta_0 + B\left(L^{1/p};\theta\right) x_{t-h}^m + \varepsilon_t \tag{1}$$

where y^q is the target variable (*quarterly* GDP growth), t indexes the lower-frequency time unit (quarterly), B is a distributed lag polynomial with parameters θ, L is the lag operator, p captures the degree of frequency mismatch ($p = 3$, for monthly/quarterly), x^m is a higher-frequency variable (*monthly* GDP or surveys), and h is the relevant forecast horizon in lower-frequency units (e.g., with one missing month, $h = 1/3$).

A key specification choice lies in the lag polynomial B. This can be expanded as:

$$B\left(L^{1/p};\theta\right) = \sum_{k=1}^{K} b\left(k;\theta\right) L^{k-1/p} \tag{2}$$

where K is the number of high-frequency lags and b their associated coefficients, which can in turn be parameterised by θ. In practice, both 'restricted' (Ghysels et al., 2007) and 'unrestricted' (U-MIDAS, Foroni et al., 2015) approaches can be used. Restricted MIDAS is more parsimonious but requires imposing a functional form that might not always be suited to the data. U-MIDAS provides greater flexibility, but at the cost of estimating a higher number of coefficients.

When debating restricted vs unrestricted MIDAS, the literature has typically emphasised the role of 'frequency mismatch', with U-MIDAS becoming more competitive as that falls. The deeper trade-off that forecasters should consider however is between i) potential accuracy gains from increased U-MIDAS coefficient flexibility, and ii) any corresponding rise in (lag) overfitting risks. In this view, the standard frequency-mismatch result is likely to reflect less U-MIDAS overfitting as the number of coefficients to be estimated drops along with frequency – for example, a U-MIDAS model with 24 weekly lags would require only 6 coefficients if recast in monthly space, while restricted MIDAS might rely on a smaller number of parameters throughout.

A corollary is that the appropriate MIDAS specification can be a function of *any* data properties affecting the above accuracy/overfitting nexus – not just frequency mismatch. A full investigation is beyond this paper's scope, but we let the logic inform out choices below.

U-MIDAS treatment of hard data

For monthly GDP data, we strictly prefer a U-MIDAS specification (Foroni et al., 2015):

$$y_t^q = \beta_0 + \sum_{k=1}^{K} b_k L^{k-1/p} y_{t-h}^m + \varepsilon_t \tag{3}$$

where y^q are first quarterly estimates, y^m are real-time-vintage lags of monthly GDP growth (Appendix B), and b are coefficients that can be estimated freely by OLS.

In the case of monthly GDP, there is a sufficiently precise relationship and sufficient data variability to make unrestricted estimation worthwhile, without overfitting. The top panel in Figure 6 shows its increasingly tighter in-sample fits as the quarter progresses. Appendix C further documents large out-of-sample gains from using U-MIDAS compared to a restricted alternative, reducing RMSE of this particular regression by as much as 30%. Our broader experience with SC-MIDAS suggests this generally holds when lower and higher-frequency versions of the same

Figure 6: Hard and soft in-sample fits, 2005-19

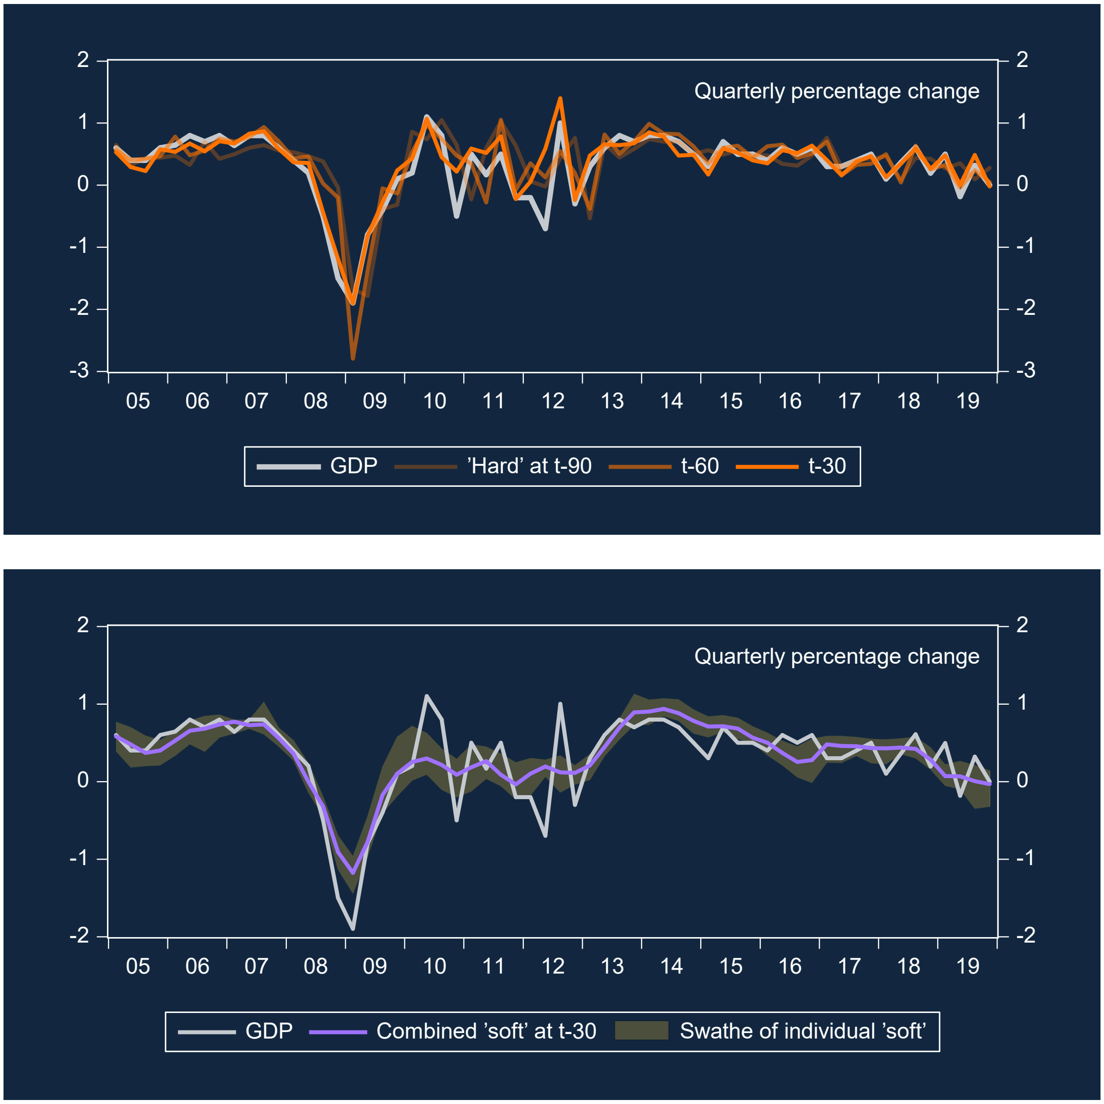

Notes: The top panel shows progressively tighter 'hard' monthly GDP fits at the 90, 60, and 30-day horizons, reflecting its 'time aggregation' properties. The bottom panel shows a swathe of individual survey predictions a 30-day horizon. A combined 'soft' measure is overlaid, which tends to smooth through official GDP volatility.

concept are used. For less precise hard indicators U-MIDAS may be less advantageous.

U-MIDAS outperforms mainly due to its ability to capture mechanical links between monthly and quarterly GDP. In a regression setting, growth in months 1, 2, and 3 of the previous and current quarters should map across to the quarterly rate with coefficients of 0, $^1/_3$, $^2/_3$, 1, $^2/_3$, and $^1/_3$, respectively. To see why, consider a *monthly growth* of 1% in month 2 ('m2') of quarter t. This raises the *level* of GDP by 1% in m2 and m3 of quarter t, but not in m1 or any of the quarter $t-1$ months, thus raising *quarterly growth* by $^2/_3$pp. To fully capture these contributions, we always

Figure 7: U-MIDAS hard coefficient estimates

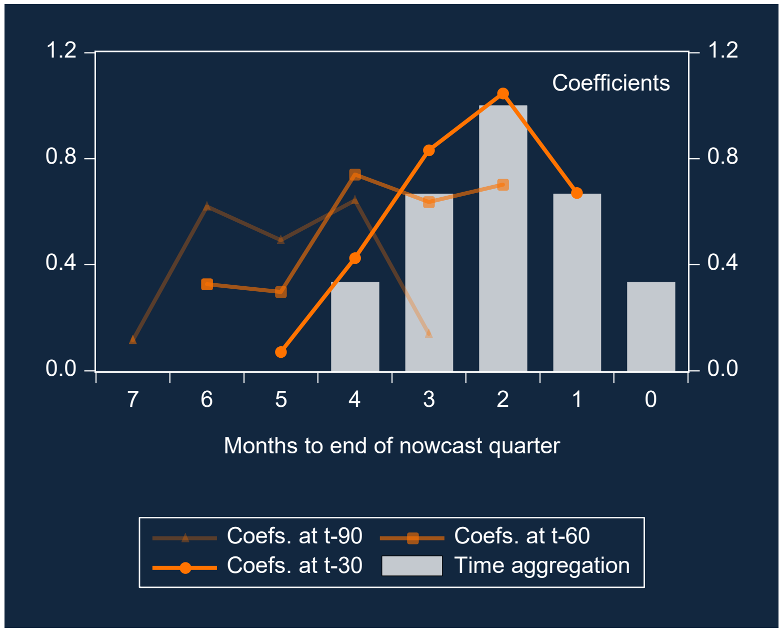

Notes: This figure shows estimated U- MIDAS coefficients through the cycle. These implicitly capture several of monthly GDP's properties, including 'time aggregation' and AR dynamics. Estimated coefficients converge toward the expected time-aggregation pattern by the end, when monthly GDP's mechanical contributions dominate.

include $K = 5$ monthly GDP lags. Figure 7 shows the corresponding U-MIDAS coefficient estimates converging steadily to the mechanical time aggregation pattern. A broader mix of factors can also be captured in those estimates, including monthly GDP's autoregressive dynamics, which are more important earlier in the quarter. Due to our use of real-time-vintage data, some of its high-frequency revision properties might also be reflected implicitly.

Almon treatment of soft indicators

For monthly surveys we prefer to use a restricted Almon approach (Almon, 1965). This involves re-writing each of the b coefficients in equation (2) as:

$$b_k = \sum_{z=0}^{Z} \theta_z k^{z-1} \tag{4}$$

where k indexes K high-frequency lags from most ($k = 1$) to least recent, $Z < K$ is the chosen lag-polynomial degree, and θ are the new parameters which in this case can also be estimated by OLS. While our setup also allows for non-linear restricted alternatives such as the *exponential* Almon or Beta lags, we have found the *linear* Almon approach to provide the best combination of parsimony and computational convenience.

Compared to monthly GDP, the accuracy/over-fitting trade-off for surveys looks less favourable, suggesting a restricted specification may be preferable. Surveys have a more diffuse link to quarterly GDP and so are less likely to benefit from additional coefficient flexibility. Meanwhile, their smoother evolution – an otherwise helpful property for anchoring the nowcast – raises risk of overfitting with multiple lags used. While appendix C shows no meaningful difference in average performance compared to U-MIDAS, Bank staff nevertheless favour Almon for practical reasons, mainly as this helps to limit coefficient volatility that can otherwise introduce spurious 'model news' across successive indicator releases.

In practice, Almon MIDAS is able to exploit survey information effectively. The shaded swathe in the bottom panel of Figure 6 encompasses their nine in-sample fits, tracking broad trends and turning points well. We include $K = 6$ lags by default to cover the typical horizons at which such surveys are informative, with polynomials of degree $Z = 3$ used to allow some 'tilt' and 'curvature' in their otherwise smooth Almon weightings. For quarterly BCC data, standard bivariate regressions with the equivalent two lags are used. We generally avoid automatic lag selections that can likewise drive unhelpful 'model news'.

3.2 Staggered combination

The staggered combination stage aggregates the set of single-indicator predictions into the full SC-MIDAS nowcast in two further steps, using a mix of RMSE-based and OLS techniques. Forecast combinations – even simple ones – can help improve forecast performance by drawing on complementary information and providing insurance against larger errors from any one model (Timmermann, 2006). The forecast error variance of an appropriately computed forecast combination can in fact be shown to never exceed that of the best individual forecasts (Bates and Granger, 1969). This idea is widely appreciated in central banks (see e.g. Bowe et al., 2023 for a recent Norges Bank application across a wide range of models and specifications), and is also leveraged *within* SC-MIDAS.

A generic forecast combination can be written as:

$$\hat{y}_{combo,t} = \sum_{i=1}^{N} w_i \hat{y}_{i,t}, \qquad \text{with} \sum_{i=1}^{N} w_i = 1 \tag{5}$$

where $\hat{y}_{combo}$ is the combined forecast, $\hat{y}_i$ are N individual forecasts, and w_i their weights.

In the two-variable case, optimal combination weights (Bates and Granger, 1969) are given by:

$$w_i = \frac{(\sigma_j^2 - \rho\sigma_i\sigma_j)}{(\sigma_i^2 + \sigma_j^2 - 2\rho\sigma_i\sigma_j)} \tag{6}$$

where σ_i^2 is model i's forecast error variance and ρ the correlation between i's and j's errors. This implies that more accurate forecasts (lower σ_i^2) tend to receive higher weight.

Optimal weights can in principle be estimated by OLS (Granger and Ramanathan, 1984), but doing so often results in the so-called 'forecast combination puzzle' – whereby finite-sample estimation error causes more sophisticated combinations to underperform simpler averaging methods (Smith and Wallis, 2009). A commonly proposed remedy is to 'ignore the covariances' ($\rho = 0$, Stock and Watson, 2004) and calculate weights based on forecasts' root mean squared errors (unbiased estimates of σ_i). This is the 'RMSE-based' combination that we use within SC-MIDAS.

In choosing the most appropriate combination method a trade-off once again emerges between the size of i) potential accuracy gains from using true optimal weights versus more naive approaches, and ii) any corresponding rise in finite-sample error risks. In this case, these two forces often move together, suggesting a simple rule of thumb. For forecasts that behave very similarly (similar σ_i^2 and high ρ) the true optimal weights will both tend to resemble equal weights *and* be harder to estimate, making simple or RMSE-based averages more attractive by default. By contrast, for forecasts that are sufficiently informative and distinct (low σ_i^2 and low ρ) optimal weights have the potential to both drive larger accuracy gains *and* be easier to identify, thus making OLS estimation more worthwhile.

RMSE-based weighting of soft signals

Because individual survey nowcasts are in practice highly collinear, we first use an RMSE-based approach to aggregate those into a single 'soft' nowcast – to then be combined with the 'hard' monthly GDP signal in a final OLS step. This produces an additional 'nowcast layer', providing a convenient summary of all available survey information.

We particularly tend to use a 'discounted, rolling-window' RMSE-based variant, with weights given by:

$$w_{i,t} = \frac{m_{i,t}^{-1}}{\sum_{j=1}^{N} m_{j,t}^{-1}}, \qquad \text{with } m_{i,t} = \sum_{v=1}^{V} \delta^v \left(y_{i,t-v} - \widehat{y}_{i,t-v}\right)^2 \tag{7}$$

where y is quarterly GDP growth, $\hat{y}_i$ are predictions from model i, V is the rolling-window length, and $\delta \in [0,1]$ is a discount factor applied to past errors.

Figure 8: RMSE-based soft combination weights, 2005-19

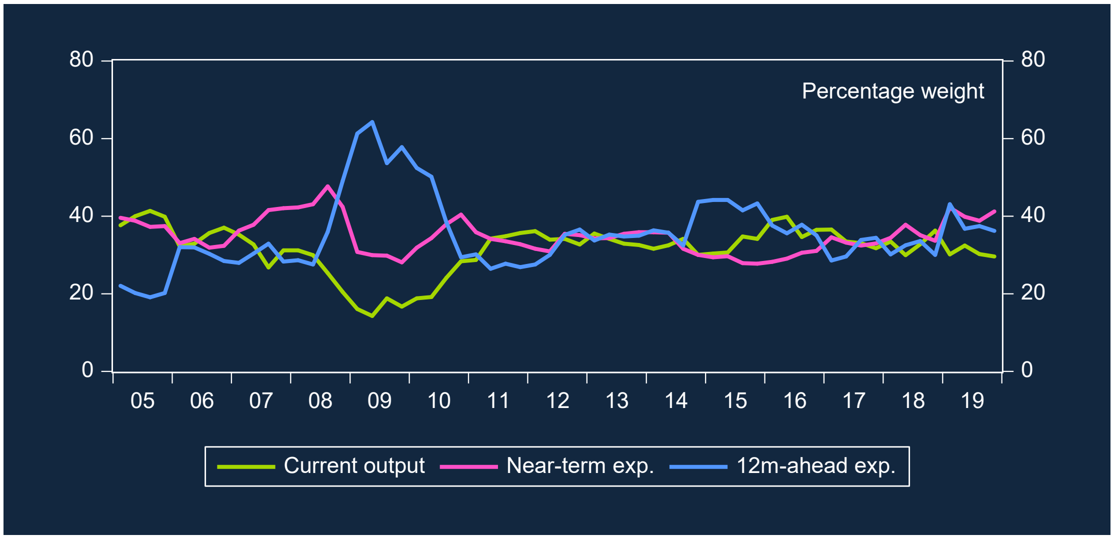

Notes: This figure shows the evolution of weights placed on different groups of surveys within the RMSE-based soft combination (not the full nowcast). Weights tend to be broadly balanced, except during the GFC, when more forward-looking measures received greater weight. Weights shown at 30 days to publication.

This adds flexibility compared to the baseline RMSE-based approach, by placing more weight on surveys that have been accurate *recently*. Helpfully compared to (6), these weights are strictly bounded between zero and one and will generally fall closer to the middle of that interval, limiting scope for overfitting. In this application, we set $\delta = 0.95$ and $V = 8$ quarters, which in our experience has helped improve performance without excessive weight variability. Appendix C provides empirical backing from a grid search exercise where the chosen settings can be seen to reduce RMSE of the combined soft measure by 3-4% compared to a simple-average benchmark.

The combined survey measure produced in this way inherits the desirable properties of its components, tracking broad trends and turning points well, whilst cancelling out series-specific noise. Figure 6 overlays that onto the swathe of survey fits. In practice, this combined measure also helpfully tend to smooth through official data noise.

The evolution of associated survey weights is also interesting to consider. Figure 8 shows those behaving reasonably neutrally over time as predicted. But they can sometimes also adjust more meaningfully – most notably during the Global Financial Crisis when greater weight was assigned to more forward-looking firm expectations. This contrasts with OLS-based weightings (not shown), which for this set of surveys would have tended to produce more extreme and unstable weights (sometimes well outside the $[0,1]$ interval). An old Bank model that used such an approach was discontinued nearly a decade ago for this reason ('weighted survey model' in Bell et al., 2014).

Figure 9: Hard vs soft optimal combination weights, 2024Q4

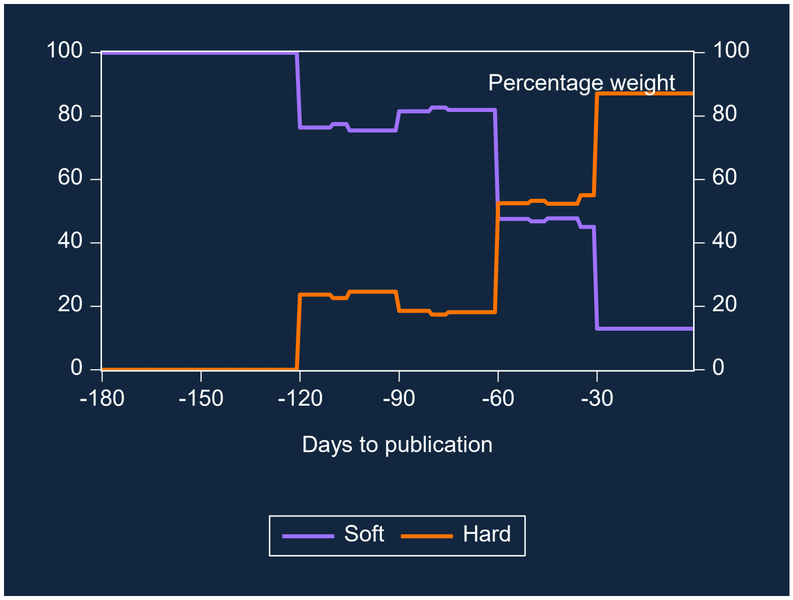

Notes: This figure shows the evolution of estimated optimal combination weights for the 'hard' vs 'soft' nowcasts over the 180-day nowcast window, for 2024 Q4. The weights swing from 100% on 'soft' at longer horizons to nearly 90% on 'hard' by the end of the window. This aggressive shifting of weights is crucial for the model's good performance through the entire nowcast window.

OLS combination of hard vs. soft signals

The 'soft' and 'hard' signals are finally blended into the 'full' SC-MIDAS nowcast in a regression setting, where in this case there is enough independent information to estimate their optimal weights by OLS (Granger and Ramanathan, 1984).

The regression is specified as follows:

$$y_t^q = \phi \widehat{y}_{soft,t}^q + (1-\phi)\widehat{y}_{hard,t}^q + \varepsilon_t, \text{ with } 0 \leq \theta_1 \leq 1 \tag{8}$$

where $\widehat{y}_{soft}^q$ and $\widehat{y}_{hard}^q$ are the 'soft' and 'hard' nowcasts, and ϕ and $(1-\phi)$ their respective combination coefficients constrained to sum to one. While we have also implemented a non-negativity constraint to guard against potential overfitting issues (as there is no obvious reason for these weights to lie outside the $[0,1]$ interval), in practice this has not tended to bind meaningfully. Greater estimation challenges may arise in applications where the hard indicator is less tightly linked to the target though. More disaggregated hard data can in principle also be used, by similarly pre-combining their signals prior to the final OLS combination. But the benefits of SC-MIDAS are likely to be largest when a single, highly informative hard indicator like monthly GDP exists.

This optimal combination regression is re-estimated continuously, enabling SC-MIDAS to re-weight soft and hard signals dynamically through the release cycle. Figure 9 plots the 180-day evolution of estimated weights for 2024Q4. Between 180 and 90 days, SC-MIDAS places only 0-20% of weight on official signals, with surveys providing the overall anchor at those horizons. In the last 90-days that picture is reversed however, with weight placed on the hard signal rising to around 50% at 60 days to publication, and 90% at 30 days.

As we will see, this flexible re-weighting of hard vs soft signals can improve accuracy compared to symmetrical pooled approaches – particularly at the final stages of the release cycle when official data become most important.

4 Performance

This section evaluates SC-MIDAS out-of-sample (OOS) and in a 'horse-race' against five other models. It also briefly addresses extensions to other UK variables, and the post-Covid period. OOS evaluation covers the 60 quarters between 2005Q1-2019Q4 on the basis of the stylised release calendar from section 2. OOS nowcasts are produced by recursively re-estimating models on real-time vintages up to quarter $t-1$ and projecting quarter t at all relevant horizons between $[-180, 0]$ days. All reported root-mean-squared errors (RMSE) exclude two quarters with large 'month 3' erratics – 2010Q4 (snow) and 2012Q2 (Jubilee bank holidays) – which, like Covid, are judged to be unrepresentative of normal economic dynamics. This lowers RMSEs slightly across the board but does not change the flavour of results. We find that SC-MIDAS' ability to exploit and dynamically reweight signals from both hard and soft data supports significant performance improvements, both through the release cycle and in comparison to other models.

4.1 OOS evaluation

Comparing the 'full' SC-MIDAS nowcasts to their 'hard' and 'soft' components helps to illustrate how the model operates and how Bank staff may draw on them in real-time. The panels in Figure 10 show those at the 90, 60 and 30-day horizons, following the releases of quarter $t-1$ GDP, and of months 1 and 2 of quarter t, respectively. A cursory look at these OOS profiles reveals a nearly identical picture to the hard and soft in-sample fits (Figure 6), attesting to the model's general stability. We next discuss OOS results at each of the three illustrated horizons.

At 90 days to publication, the full model nowcast typically tracks soft signals very closely. This echoes the previous discussion of SC-MIDAS' weighting of hard vs soft signals, where surveys were shown to receive about 80% of weight at this horizon (Figure 9). Anchoring the nowcast

Figure 10: SC-MIDAS nowcasts at 90, 60 and 30-day horizons, 2005-19

Notes: This figure shows OOS nowcasts results from the 'full' SC-MIDAS model and its 'hard' and 'soft' components, plotted against first estimates of quarterly GDP growth. SC-MIDAS performs well through the release cycle, becoming notably more accurate as monthly GDP information becomes available.

Figure 11: SC-MIDAS RMSEs through the release cycle

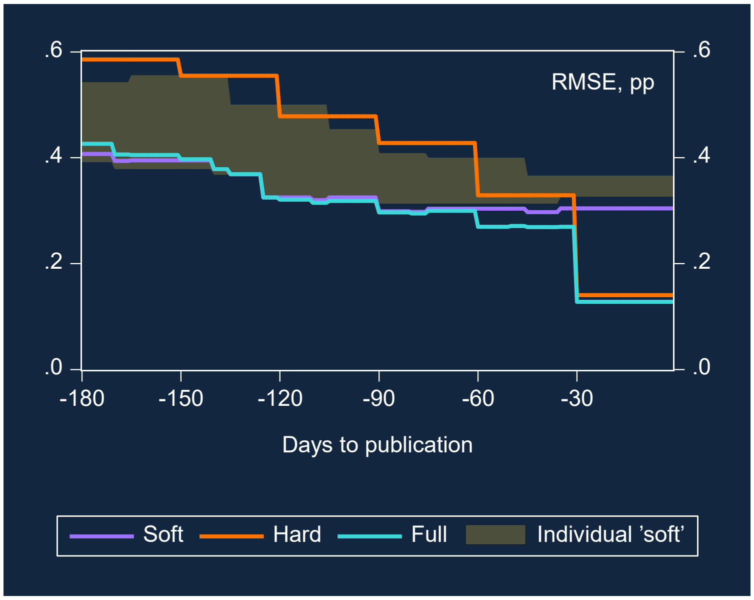

Notes: This figure shows the evolution of 2005-19 RMSEs for the 'full' SC-MIDAS nowcast and its 'hard' and 'soft' components through the 180-day nowcast window. The full model generally performs at least as well as the components. The shaded band corresponds to the swathe of individual survey RMSEs.

to timelier surveys enables SC-MIDAS to track broad growth trends and turning points well, while avoiding the more erratic behaviour of the hard signal during certain periods, for example 2008-12.

At 60 days to publication, the nowcast begins to resemble a simple average of soft and hard signals, as more weight is placed on the latter following 'month 1' GDP. This strikes a reasonable balance between the still informative survey anchor and incoming hard information, which can often begin to foreshadow GDP's quarterly swings at this stage. In several instances however, putting full stock on monthly GDP can still lead to outsized errors. This suggests an important role for staff judgement in assessing the reliability of monthly GDP's latest signals on a case-by-case basis. For example, if staff have reason to believe that a recent outturn is affected by a sharp one-off distortion, they may wish to place *less* weight on hard signals than suggested by the combination. In the absence of any such distortions, they may at times wish to place *more* weight on that instead.

Finally, at 30 days to publication, following the release of 'month 2' GDP, the full SC-MIDAS nowcast becomes virtually indistinguishable from the hard signal. With five out of six inputs to the quarterly rate available, time aggregation naturally becomes the dominant driver of the nowcast. As discussed in 3 SC-MIDAS is designed to exploit this fully, through its deliberate use of a flexible U-MIDAS specification and optimal re-weighting of hard and soft signals through the quarter – which at this stage allows as much as 90% of the weight to be placed on official information. The resulting nowcasts are extremely accurate, leaving limited room for staff judgements unless exceptional 'month 3' erratics (such as the aforementioned 2010Q4 snow or 2012Q2 Jubilee celebrations) are known to have occurred.

Average model performance is summarised in Figure 11, which plots the evolution of each measure's RMSE. The clear 'staircase' pattern reflects steady increases in accuracy as SC-MIDAS effectively extracts and combines signals from the incoming data flow. Its headline RMSE drops dramatically from around 0.4 percentage points (pp) at 180 days, to 0.3pp at 90 days, and just over 0.1pp at 30 days. Within the breakdown, the full SC-MIDAS measure always performs at least as well as its hard and soft components, echoing the classic result from the forecast combination literature (Bates and Granger, 1969). The same is true of the combined *soft* measure itself, whose RMSE generally lies below the swathe of individual surveys.

4.2 Horse-race

We also evaluate SC-MIDAS in a 'horse-race' against five other models, including quarterly and monthly autoregressive benchmarks (referred to as 'Q-AR' and 'M-AR', respectively), a monthly Bridge-style model ('Bridge'), a Dynamic Factor Model ('DFM'), and a standard Pooled-MIDAS model ('P-MIDAS'). The trio of Bridge, DFM, and P-MIDAS is similar, though not exactly identical, to the Bank's old three-model system (Anesti et al., 2017). In what follows, each of the models is evaluated OOS on the same real-time information set as SC-MIDAS, such that any performance differences can be attributed to methods. We next provide a brief description of each of the five competitors, before turning to the horse-race results.

Quarterly AR. The Q-AR benchmark employs a canonical autoregressive specification:

$$y_t^q = \beta_0 + \beta_1 y_{t-1}^q + \varepsilon_t \tag{9}$$

where y^q represents quarterly GDP growth, and a single lag is used. No higher-frequency or soft information is included.

Monthly AR. The M-AR model uses a monthly AR(2) specification:

$$y_t^m = \beta_0 + \sum_{i=1}^{N} \beta_i y_{t-i}^m + \varepsilon_t \tag{10}$$

where y^m represents monthly GDP growth, and where the quarterly nowcasts of interest can be obtained by conventional time-aggregation of the model's iterative forecasts. The choice of an AR(2) is pre-informed by out-of-sample testing of different lag-lengths. Once again, no soft information is included.

Dynamic factor model. The DFM used here is a version of Bok et al. (2018) at the New York Fed, adapted by Bank colleagues to be recursively estimated and evaluated on real-time UK data. The model is summarised by the following measurement equation:

$$x_{i,t} = \lambda_i f_t + e_{i,t} \tag{11}$$

where x is a set of variables capturing the exact same information set as SC-MIDAS, f is the 'common factor', λ are factor loadings, and e_i is an idiosyncratic component capturing variable-specific movements for each i. Transition equations for the common factor and idiosyncratic components are assumed to follow Gaussian AR(1) processes. The system is estimated using the Kalman smoother and the expectation-maximisation (EM) algorithm (Bańbura and Modugno, 2014), which can handle mixed frequencies and missing observations.

Bridge. The Bridge model is based on an approach outlined by Bank colleagues in Bell et al. (2014). This begins by estimating a set of monthly AR(2) regressions for each soft indicator:

$$x_t^m = \gamma_0 + \sum_{i=1}^{N} \gamma_i x_{i,t-i}^m + \varepsilon_t \tag{12}$$

where x_i^m represents monthly observations of soft indicator i. The forecasts from this are then plugged into a corresponding set of ARDL regressions of the form:

$$y_t^m = \beta_0 + \sum_{i=1}^{N} \beta_i y_{t-i}^m + \beta_{N+1} \widehat{x}_{i,t}^m + \varepsilon_t \tag{13}$$

where y^m is monthly GDP growth, $\widehat{x}^m$ is the forecast of monthly indicator i, and quarterly GDP nowcasts are once again obtained by time-aggregation. While this differs slightly from typical bridge models, where equation (13) would normally be set at lower frequency (Baffigi et al., 2004), it is similar in spirit and enables us to better exploit monthly GDP information. The overall Bridge nowcast is obtained by pooling the resulting set of quarterly predictions using the RMSE-based method from section 3.

Pooled-MIDAS. The P-MIDAS approach corresponds closely to SC-MIDAS' predecessor at the Bank, the more standard pooled-MIDAS model documented in Anesti, Hayes, Moreira and Tasker (2017). Similar to SC-MIDAS, this begins with a set of MIDAS regressions for each indicator:

$$y_t^q = \beta_0 + B\left(L^{1/p};\theta\right) x_{i,t-h}^m + \varepsilon_t \tag{14}$$

Figure 12: Horse-race RMSEs through the release cycle

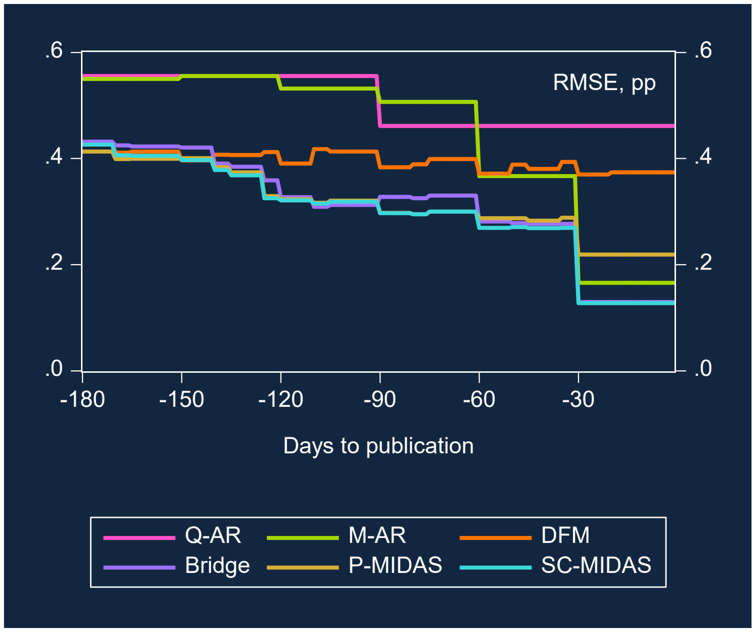

Notes: This figure shows the evolution of 2005-19 RMSEs for SC-MIDAS and the five competitors through the 180-day nowcast window. SC-MIDAS generally has the lowest RMSE, matching both P-MIDAS and the DFM's ability to exploit survey data early in the window, and mimicking Bridge's 'time aggregation' of monthly GDP contributions towards the end.

where y^q is quarterly GDP growth and x_i^m the relevant predictor. For the cleanest comparison, the exact same individual MIDAS specifications from section 3 are used. The only difference therefore lies in the combination stage, where P-MIDAS pools all predictions symmetrically through a single RMSE-based step, compared to SC-MIDAS' two-step combination.

Horse-race results

We find that SC-MIDAS generally outperforms the five horse-race competitors. Figure 12 plots the evolution of their RMSEs over the 180-day nowcast window, with raw OOS nowcasts also included for completeness in Appendix D. Table 2 complements those with a set of *relative* RMSE statistics and Diebold-Mariano (DM) tests (Diebold and Mariano, 2002) of comparative predictive accuracy. In Figure 12, the SC-MIDAS line broadly traces out the 'lower envelope' of RMSEs – generally matching or outperforming that of the single best competitor at each stage of the release cycle.

At the start of the 180-day window, differences between models that include soft information are not significant. These in turn beat the Q-AR and M-AR benchmarks clearly. As the release cycle

Table 2: SC-MIDAS vs. other models, relative RMSEs and Diebold-Mariano (DM) tests

Horizon:	**180**	**150**	**120**	**90**	**60**	**30**	**Average**
Q-AR	0.77	0.68	0.58	0.64	0.59	0.28	0.59
(DM p-value)	*(0.003***)*	*(0.002***)*	*(0.013**)*	*(0.002***)*	*(0.001***)*	*(0.001***)*	–
M-AR	0.77	0.68	0.60	0.59	0.74	0.77	0.67
(DM p-value)	*(0.011**)*	*(0.018**)*	*(0.024**)*	*(0.001***)*	*(0.038**)*	*(0.044**)*	–
DFM	1.00	0.93	0.82	0.77	0.70	0.34	0.76
(DM p-value)	*(0.620)*	*(0.318)*	*(0.007***)*	*(0.007***)*	*(0.006***)*	*(0.000***)*	–
Bridge	0.98	0.96	0.95	0.89	0.97	0.98	0.96
(DM p-value)	*(0.293)*	*(0.107)*	*(0.143)*	*(0.067*)*	*(0.3723)*	*(0.395)*	–
P-MIDAS	1.03	0.99	0.99	1.00	0.94	0.58	0.94
(DM p-value)	*(0.939)*	*(0.104)*	*(0.331)*	*(0.434)*	*(0.0842*)*	*(0.005***)*	–

Notes: This table shows relative RMSES of SC-MIDAS versus each of the horse-race competitors. Values below 1.0 denote improvement. SC-MIDAS generally outperforms through the entire, or most of the nowcast window, reducing RMSE by up to 66%, 42% and 11% compared to DFM, P-MIDAS and Bridge. Green denotes a statistically significant difference according to a standard DM test (***, **, * denote significance at 1%, 5% and 10% levels).

progresses, incoming survey data (and to a lesser extent monthly GDP) help to reduce RMSEs of SC-MIDAS, P-MIDAS and Bridge from over 0.4pp to around 0.3pp at the 90-day mark. By that point, DM tests detect statistically significant improvements over both DFM and Bridge, with SC-MIDAS reducing RMSE by 23% and 11% respectively, but not P-MIDAS, whose symmetrical pooling of individual forecasts still approximates SC-MIDAS' optimal weighting of hard vs soft group-level signals reasonably well.

In the final 90 days, the arrival of monthly GDP data gives SC-MIDAS more of an edge. Following the publication of months 1 and 2, SC-MIDAS' RMSE drops sharply from 0.3pp at 90 days to just over 0.1pp at the 30-day horizon. Performance improvements over the DFM and P-MIDAS then become highly statistically significant according to DM tests, with SC-MIDAS cutting RMSE by as much as 66% and 42% relative to each of those. The only model that is able to match SC-MIDAS' effective use of hard data at this stage is Bridge, which has conventional time aggregation built in by design.

The relatively poor performance of the DFM echoes past Bank research which has also found those to perform relatively modestly in UK GDP applications (Anesti et al., 2017, 2022). However, with DFMs designed to handle larger datasets, performance in this case may have been hampered by the relatively small number of input series. The Bridge model on the other hand appears considerably more competitive than previous Bank research had found (Anesti et al., 2017). These differences probably reflect our use of full real-time vintages, and the fact that an additional combination step was built in to also help that leverage signals from a broader range of soft indicators.

By matching P-MIDAS and the DFM's ability to exploit surveys at the start of the nowcast window, *and* mimicking Bridge's conventional time-aggregation of monthly GDP toward the end, SC-MIDAS has effectively combined the best performance features of the Bank's old three-model system within a single model.

4.3 Other applications

Bank staff have also applied SC-MIDAS successfully to a range of other UK variables where the lower-frequency target variable is sampled at higher frequency. The growing list includes sectoral output data for services, manufacturing, and construction; retail sales; expenditure components such as exports and imports; key labour market variables such as employment and pay; and prices data such as CPI and house prices inflation. We intend to return to some of those in future Macro Technical Papers.

Our positive experience so far suggests SC-MIDAS is likely be applicable more widely – not just across variables but potentially also countries. As a brief preview of other existing Bank applications, figure 13 shows equivalent 2005-19 out-of-sample performance metrics for UK employment and (quarterly, seasonally adjusted, and excluding volatile energy) CPI inflation, where SC-MIDAS tends to behave remarkably similarly. A more detailed discussion of labour market results, for which the method is also well established, has already been offered in Daniell and Moreira (2023).

4.4 Covid and beyond

Covid posed significant challenges for nowcast models, severely disrupting standard dynamics and relationships on which they rely. Traditional soft indicators in particular broke down, as their qualitative 'net balance' nature ('up', 'down', 'same') made them ill-suited to capture the full scale of pandemic swings. At the Bank, we paused our use of SC-MIDAS and other nowcast models as a result, turning instead to a range of unconventional high-frequency indicators to help us monitor the economy in real time (Bank of England, 2020).

SC-MIDAS has since been brought back online, with all Covid quarters between 2020Q1 and 2021Q4 excluded through a set of period dummies. This is a straightforward and effective way of dealing with Covid-related estimation distortions, at least when the primary aim is to produce

Figure 13: SC-MIDAS RMSEs in different applications

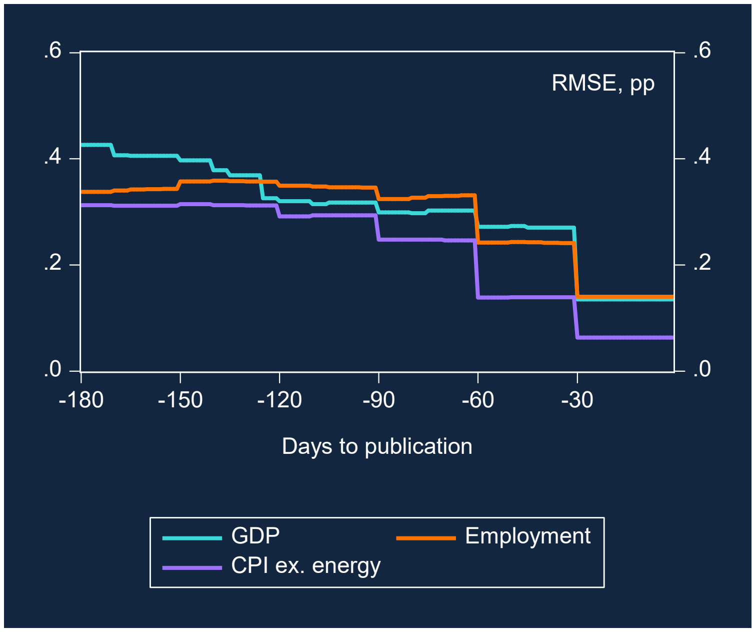

Notes: This figure shows 2005-19 SC-MIDAS RMSEs through the 180-day nowcast window for three applications currently in use by Bank staff, showing results for employment and CPI alongside GDP.

Figure 14: SC-MIDAS nowcasts since 2020

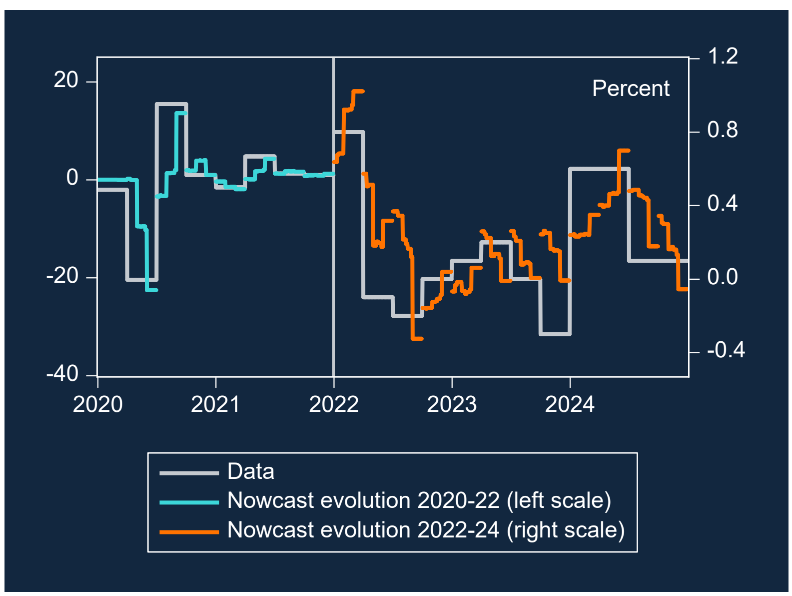

Notes: This figure shows the evolution of SC-MIDAS nowcasts through each of the quarter between 2020Q1 and 2024 Q4. In 2020 and 2021 the model is unsurprisingly less able to anticipate the extent of extreme Covid swings at longer horizons, but continues to do well once enough monthly GDP information gets published toward the end of the quarter. Performance since then appears broadly in line with pre-Covid norms.

point estimates. For inference purposes, the literature has highlighted that greater care is needed (Lenza and Primiceri, 2022).

Evidence so far suggests SC-MIDAS' performance remains intact post-Covid. Figure 14 shows the evolution of OOS SC-MIDAS nowcasts over 2020Q1-2024Q4. As before, these tend to become notably more accurate as the quarter progresses, with post-Covid RMSEs dropping from around 0.35pp at 90 days to just 0.2pp at 30 days – a little higher than pre-pandemic but in the same ballpark (Figure 11). Also interesting in retrospect is the fact that even in 2020-2021 (notice the different scale in Figure 14), SC-MIDAS would have been able to deliver fairly accurate late-stage nowcasts, thanks to its effective use of monthly GDP data. For example, in 2020Q1, the model's final OOS prediction would have been -22.5% vs a first estimate of -20.4%. In 2020Q2 it would have pointed to +13.6% vs a first estimate of +15.5%.

5 Conclusion

This paper has introduced a Staggered-Combination MIDAS (SC-MIDAS) approach, developed and used by Bank of England staff to inform nowcasts of key macroeconomic variables and brief the Monetary Policy Committee. While SC-MIDAS is built entirely upon standard techniques, its bespoke design – structured explicitly around 'hard' and 'soft' data as distinct informational groupings with signals to be combined – differentiates it from more conventional pooled-MIDAS approaches. Out-of-sample testing showed that this can result in significant performance gains, both through the release cycle and relative to other models. Our positive experience with other UK applications suggests SC-MIDAS is applicable more widely.

Looking ahead, the Bank is considering how to more effectively extract and combine insights from multiple models for various policy-relevant purposes, including short-term forecasting. We will provide further updates on staff's preferred approaches as we continue to improve our toolkit in response to the Bernanke review of forecasting for monetary policy making and communication at the Bank of England (Bernanke, 2024; Bank of England, 2024; Lombardelli, 2024).

Appendix

A Raw data inputs

Figure 15: Monthly GDP growth vintages

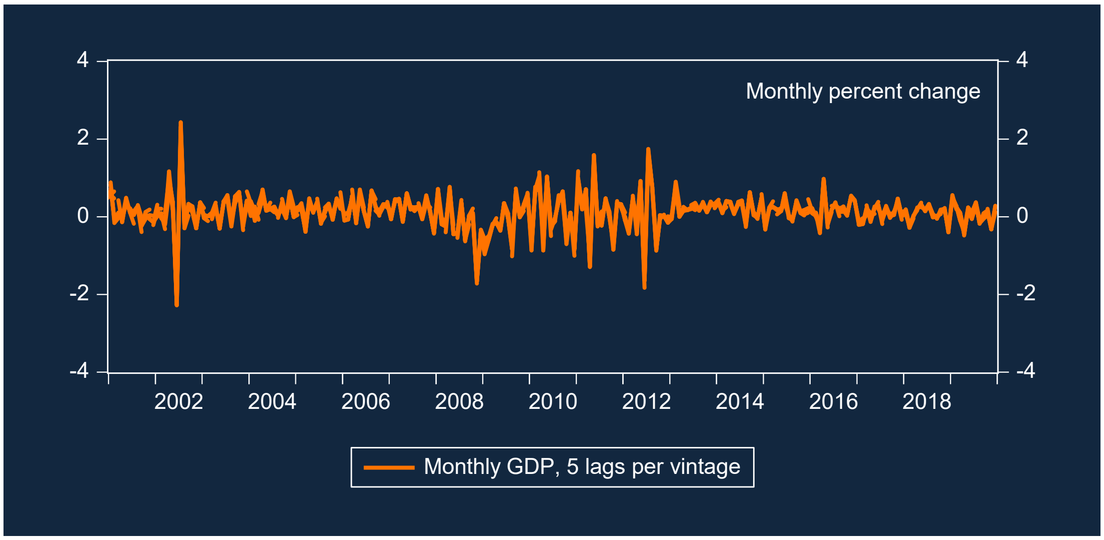

Figure 16: Survey indicators, standardised

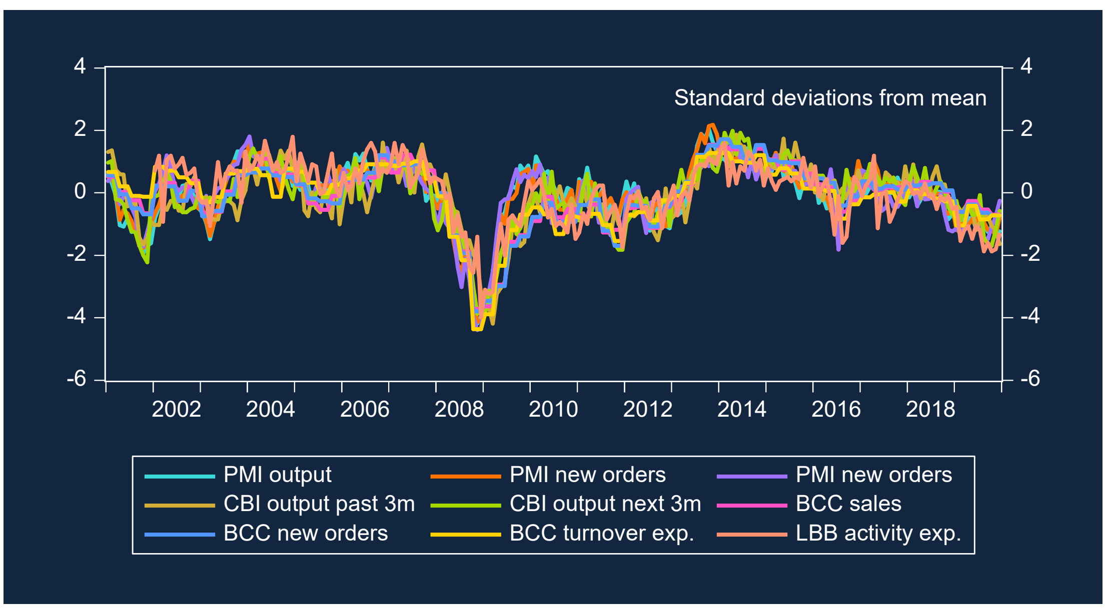

B Using real-time vintages

This appendix explains how real-time-vintage GDP data can be prepared for estimation. In our chosen application, SC-MIDAS is set up to target 'first quarterly estimates' of GDP growth. For consistency, it is important to also use full 'real-time-vintage' (Koenig et al., 2003) monthly GDP data on the right-hand-side of the model, so that can be evaluated based on actual information that would have been available to forecasters at the time. Although here we focus on GDP, the same principles apply to any variables subject to revision.

As explained in section 3, SC-MIDAS uses an 'unrestricted MIDAS' specification (U-MIDAS Foroni et al., 2015) to extract signals from monthly GDP data. To enable full real-time estimation, monthly GDP vintages need to be rearranged in a particular way. This involves creating *quarterly* time series for each of the relevant *monthly* lags. For example, if the forecaster wishes to predict 2019Q4 on the basis of data to November (month 2), they would need to create quarterly series of all 'month 2s', 'month 1s', etc., taken from the equivalent stage in the release cycle. Even though mixed-frequency data are used, the resulting regression is effectively cast in quarterly space and can be estimated by OLS. Figure 17 illustrates the process for $K = 5$ monthly GDP lags.

Figure 17: Rearranging real-time vintages for estimation

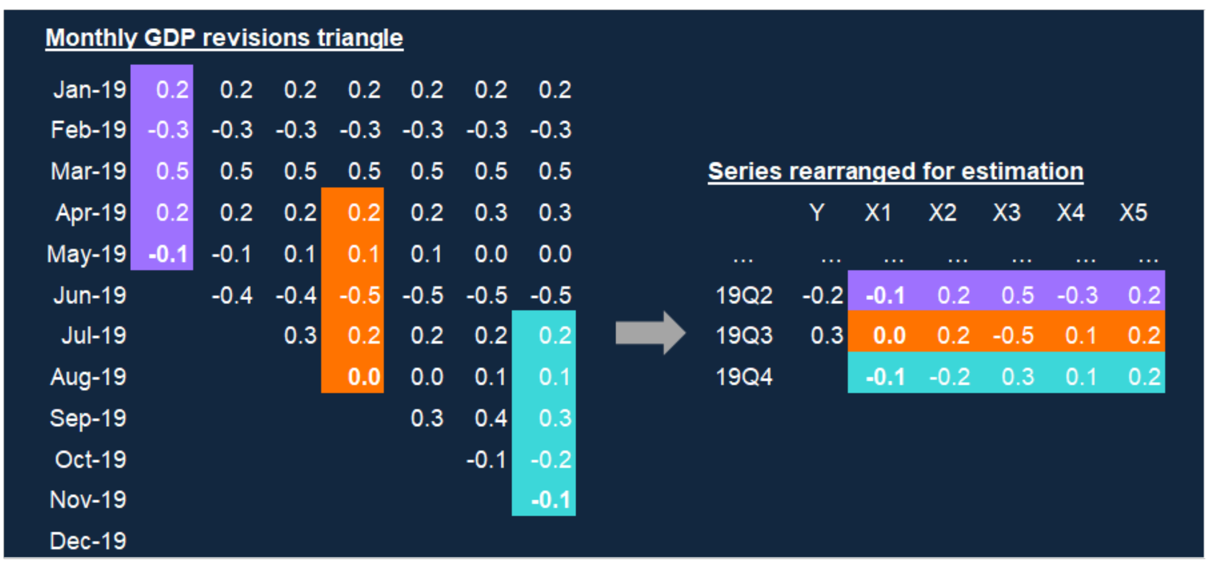

C Sensitivity analysis

This appendix documents sensitivity analysis around a few SC-MIDAS specification choices introduced in section 3, first exploring the performance of 'restricted' vs 'unrestricted' MIDAS specifications in extracting signals from hard and soft data, then investigating the role of window-length and discount-factor settings within the RMSE-based combination step.

Unrestricted vs Almon MIDAS

SC-MIDAS uses unrestricted MIDAS (U-MIDAS, Foroni et al., 2015) and restricted Almon MIDAS regressions (Almon, 1965; Ghysels et al., 2004) to extract signals from hard and soft indicators, respectively. Figure 18 explores how nowcast accuracy of each of those groupings changes if the other specification is used instead. Results are based on out-of-sample testing over 2005-19 and expressed in 'relative RMSEs', where a number below 1 signifies U-MIDAS outperforms. In the case of hard data, U-MIDAS reduces RMSE by as much as 30% compared to Almon MIDAS (relative RMSE of 0.7), as U-MIDAS better captures time-aggregation links toward the end of the nowcast window. At the start of the window U-MIDAS underperforms however, as higher coefficient numbers become more of a liability at longer horizons. For soft data, average performance is broadly comparable across U-MIDAS and Almon specifications.

Figure 18: U-MIDAS vs Almon, relative RMSEs through the release cycle

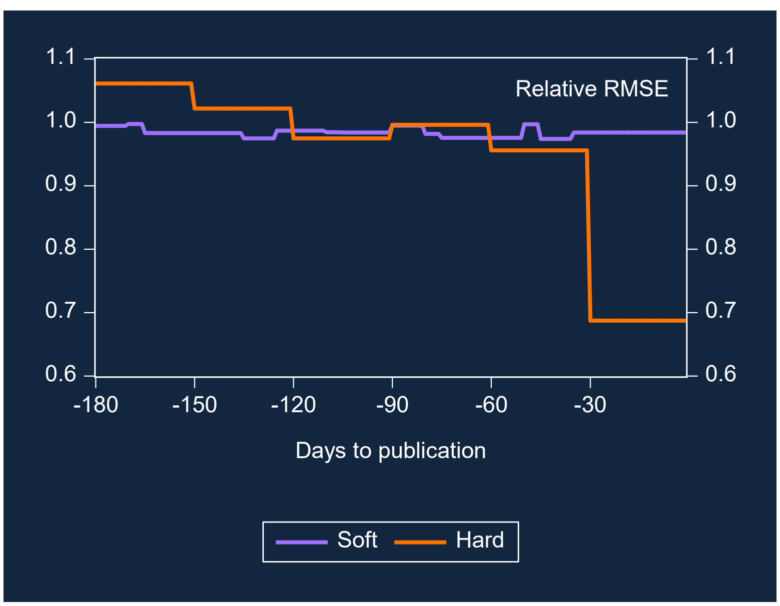

RMSE-based combination

SC-MIDAS' staggered combination begins with an RMSE-based weighting (Stock and Watson, 2004) of individual soft signals into a single, combined 'soft' nowcast. Settings used in this paper, where those weights are calculated on the basis of a rolling window of $V = 8$ quarters (2 years) and applying a discount factor of $\delta = 0.95$ to past errors, are broadly similar to those typically used by Bank staff in practice, although note these can be subject to review. This approach intuitively assigns more weight to more recently accurate indicators, with older errors penalised less until they fall out of the defined window.

Figure 19 presents results from a grid-search exploring accuracy of different window/discount pairs, showing relative RMSEs of soft nowcasts combined in this way compared to a simple average (equal weights) benchmark. Results shown are based on 2005-19 out-of-sample testing at the 90-day horizon, when surveys are still the dominant force in SC-MIDAS' nowcast. The overarching lesson is that moderate degrees of weight adaptability tend to be helpful in this case, reducing RMSE by up to 4% compared to the simple-average benchmark. This adaptability can either be achieved through use of short rolling windows up to 2-3 years, or through more aggressive discounting. Excessive discounting becomes actively harmful however, as RMSEs begin to rise again for lower factors. The used pairing of $V = 2$ and $\delta = 0.95$ sits broadly on the blue region, where performance gains are largest.

Figure 19: RMSE-based soft combinations vs simple average, relative RMSE

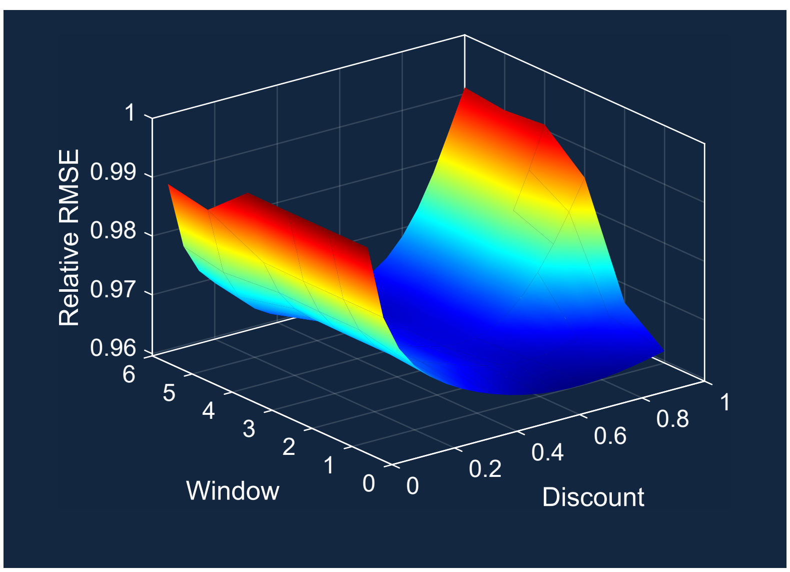

D Horse-race nowcasts

Figure 20: Horse-race nowcasts through the release cycle

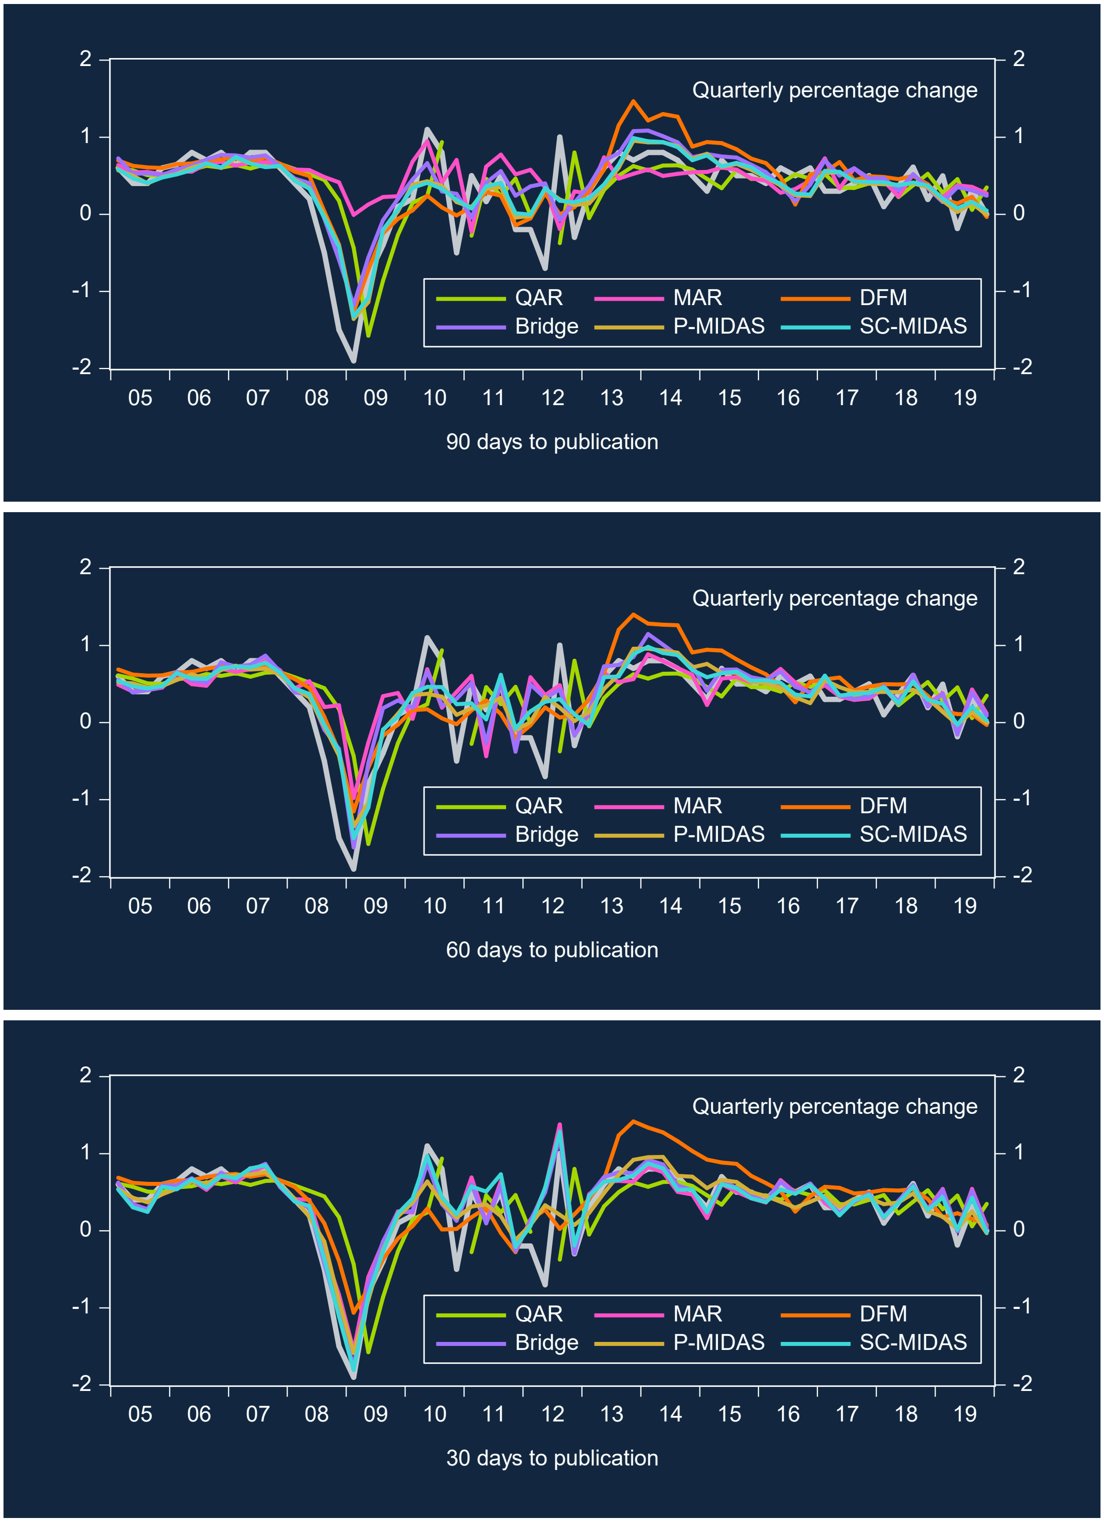

Notes: This figure shows snapshots of OOS nowcasts from the six models in the 'horse-race' from Section 4.

References

Albuquerque, Daniel, Jenny Chan, Derrick Kanngiesser, David Latto, Simon Lloyd, Sumer Singh, and Jan Zacek (2025) "Decompositions, forecasts and scenarios from an estimated DSGE model for the UK economy," *Bank of England Macro Technical Paper*.

Almon, Shirley (1965) "The distributed lag between capital appropriations and expenditures," *Econometrica*, pp. 178–196.

Anesti, Nikoleta, Ana Beatriz Galvão, and Silvia Miranda-Agrippino (2022) "Uncertain Kingdom: Nowcasting Gross Domestic Product and its revisions," *Journal of Applied Econometrics*, Vol. 37.

Anesti, Nikoleta, Simon Hayes, Andre Moreira, and James Tasker (2017) "Peering into the Present: the Bank's Approach to GDP Nowcasting," *Bank of England Quarterly Bulletin 2017 Q2*.

Angelini, Elena, Gonzalo Camba-Mendez, Domenico Giannone, Lucrezia Reichlin, and Gerhard Rünstler (2011) "Short-term forecasts of euro area GDP growth," *The Econometrics Journal*, Vol. 14, pp. C25–C44.

Baffigi, Alberto, Roberto Golinelli, and Giuseppe Parigi (2004) "Bridge models to forecast the euro area GDP," *International Journal of Forecasting*, Vol. 20, pp. 447–460.

Bańbura, Marta, Domenico Giannone, Michele Modugno, and Lucrezia Reichlin (2013) "Nowcasting and the real-time data flow," Vol. 2: Elsevier, Chap. Chapter 4, pp. 195–237.

Bańbura, Marta and Michele Modugno (2014) "Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data," *Journal of applied econometrics*, Vol. 29, pp. 133–160.

Bańbura, Marta and Gerhard Rünstler (2011) "A look into the factor model black box: publication lags and the role of hard and soft data in forecasting GDP," *International Journal of Forecasting*, Vol. 27, pp. 333–346.

Bank of England (2018) "Implications of changes to the ONS GDP publication timetable," *Inflation Report, August 2018*, p. 11.

——— (2020) "How are we monitoring the economy during the Covid-19 pandemic?" *Bank Overground blog*.

——— (2024) "Response of the Bank of England to the Bernanke review of forecasting for monetary policy making and communication at the Bank of England."

——— (2025) "Current economic conditions," *Monetary Policy Report, May 2025*, p. 35.

Bates, John M. and Clive W. J. Granger (1969) "The combination of forecasts," *Journal of the Operational Research Society*, Vol. 20, pp. 451–468.

Bell, Venetia, Lai Wah Co, Sophie Stone, and Gavin Wallis (2014) "Nowcasting uk gdp growth," *Bank of England Quarterly Bulletin 2014 Q1*.

Bernanke, Ben (2024) "Forecasting for monetary policy making and communication at the Bank of England: a review,"Technical report.

Bok, Brandyn, Daniele Caratelli, Domenico Giannone, Argia M. Sbordone, and Andrea Tambalotti (2018) "Macroeconomic nowcasting and forecasting with big data," *Annual Review of Economics*, Vol. 10, pp. 615–643.

Bowe, Frida, Inga Nielsen Friis, Atle Loneland, Erlend Salvesen Njølstad, Sara Skjeggestad Meyer, Kenneth Sæterhagen Paulsen, and Ørjan Robstad (2023) "A SMARTer way to forecast," staff memo, Norges Bank.

Broadbent, Ben (2023) "Signal versus noise," December, Speech given at the London Business School.

Burgess, Stephen, Emilio Fernandez-Corugedo, Charlotta Groth, Richard Harrison, Francesca Monti, Konstantinos Theodoridis, and Matt Waldron (2013) "The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models," Bank of England working papers 471, Bank of England.

Clements, Michael P and Ana Beatriz Galvão (2009) "Forecasting US output growth using leading indicators: An appraisal using MIDAS models," *Journal of Applied Econometrics*, Vol. 24, pp. 1187–1206.

Cunningham, Alastair, Jana Eklund, Chris Jeffery, George Kapetanios, and Vincent Labhard (2012) "A state space approach to extracting the signal from uncertain data," *Journal of Business & Economic Statistics*, Vol. 30, pp. 173–180.

Daniell, Harvey and Andre Moreira (2023) "Forecasting near-term trends in the labour market," *Bank Underground blog*.

Diebold, Francis X. and Robert .S Mariano (2002) "Comparing predictive accuracy," *Journal of Business & economic statistics*, Vol. 20, pp. 134–144.

Foroni, Claudia, Massimiliano Marcellino, and Christian Schumacher (2015) "Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, Vol. 178, pp. 57–82.

Ghysels, Eric and Massimiliano Marcellino (2018) *Applied Economic Forecasting Using Time Series Methods*: Oxford University Press.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov (2004) "The MIDAS Touch: Mixed Data Sampling Regression Models," cirano working papers, CIRANO.

Ghysels, Eric, Arthur Sinko, and Rossen Valkanov (2007) "MIDAS regressions: Further results and new directions," *Econometric Reviews*, Vol. 26, pp. 53–90.

Giannone, Domenico, Lucrezia Reichlin, and David Small (2008) "Nowcasting: The real-time informational content of macroeconomic data," *Journal of Monetary Economics*, Vol. 55, pp. 665–676.

Granger, Clive W. J. and Ramu Ramanathan (1984) "Improved methods of combining forecasts," *Journal of Forecasting*, Vol. 3, pp. 197–204.

Higgins, Patrick C. (2014) "GDPNow: A Model for GDP Nowcasting," FRB Atlanta Working Paper 2014-7, Federal Reserve Bank of Atlanta.

Koenig, Evan F., Sheila Dolmas, and Jeremy Piger (2003) "The use and abuse of real-time data in economic forecasting," *Review of Economics and Statistics*, Vol. 85, pp. 618–628.

Kohns, David and Galina Potjagailo (2023) "Flexible Bayesian MIDAS: time-variation, group-shrinkage and sparsity," Bank of England working papers 1025, Bank of England.

Lenza, Michele and Giorgio E Primiceri (2022) "How to estimate a vector autoregression after March 2020," *Journal of Applied Econometrics*, Vol. 37, pp. 688–699.

Lombardelli, Clare (2024) "Managing the present, shaping the future," November, Speech given at the 3rd Bank of England Watchers Conference.

Mann, Catherine L. (2025) "The February 2025 Monetary Policy Report - and my vote," February, Speech given at Leeds Beckett University.

Mantoan, Giulia and Jessica Verlander (2025) "Beyond point estimates: quantifying risk around the near-term UK GDP forecast using a new quantile-MIDAS model," *Bank Underground blog*.

Mitchell, James (2009) "Where are we now? The UK recession and nowcasting GDP growth using statistical models," *National Institute Economic Review*, Vol. 209, pp. 60–69.

ONS (2016) "National Accounts methodology."

——— (2018) "Introducing a new publishing model for GDP."

Ramsden, Dave (2025) "Surveys, forecasts and scenarios: setting UK monetary policy under uncertainty," February, Keynote speech at Stellenbosch University, South Africa.

Robinson, Lee (2016) "Rewriting history: understanding revisions to UK GDP," *Bank Underground blog*.

Smith, Jeremy and Kenneth F. Wallis (2009) "A simple explanation of the forecast combination puzzle," *Oxford Bulletin of Economics and Statistics*, Vol. 71, pp. 331–355.

Stock, James H. and Mark W. Watson (2004) "Combination forecasts of output growth in a seven-country data set," *Journal of Forecasting*, Vol. 23, pp. 405–430.

Timmermann, Allan (2006) "Forecast Combinations," Vol. 1: Elsevier, 1st edition, Chap. 4, pp. 135–196.