Office for
Budget Responsibility

Working paper No.17

Evaluating forecast uncertainty with stochastic simulations

Daniel Steel

December 2021

Evaluating forecast uncertainty with stochastic simulations

Daniel Steel

Office for Budget Responsibility

Abstract

Currently at the OBR we use a variety of approaches to illustrate uncertainty around our central forecasts. This paper adds stochastic simulations to our toolkit. It sets out our approach to generate forecast distributions for variables of interest, particularly debt-to-GDP, which involves a vector autoregression (VAR) model to which we append a debt accumulation identity. This approach has several advantages over our current historical forecast error approach including allowing us to capture a longer and therefore potentially more representative history of shocks affecting the UK economy. It also allows us to assess the probability of meeting a wider variety of fiscal rules, both individually and jointly. We plan to use this as our primary approach in the future for generating fan charts and assessing fiscal rules.

I would like to thank the members of the Budget Responsibility Committee Charlie Bean, Andy King, and Richard Hughes, and Office for Budget Responsibility colleagues Edward Robson, Edward Egan, Rosie Colthorpe, Tom Wickstead, Rohan Sheth, Harriet Brown and Kate Hall-Strutt for their comments and invaluable contributions to this paper.

1 Introduction and summary

1.1 At least twice a year the OBR publishes a set of central (strictly speaking median) forecasts for the economy and public finances in our *Economic and fiscal outlooks* (*EFO*). These central forecasts invariably turn out to be wrong to some degree – in the case of shocks like the pandemic, to a substantial degree – so we regularly evaluate the sources of those errors in our annual *Forecast evaluation reports*. But rather than focus solely on our central forecasts, it can be more useful to think of a distribution of possible future outcomes with the eventual outturn being just one possible realisation.

1.2 Currently at the OBR we use a variety of approaches to illustrate the uncertainty around our central forecast, including fan charts based on historical errors. In this paper we describe the use of stochastic simulations to generate forecast distributions for variables of interest, notably public sector net debt (PSND), which is a natural outgrowth of this approach. Using stochastic simulations instead for fan charts, a method already used by organisations such as the IMF, has several advantages which include:

- allowing us to capture a longer and potentially more representative history of shocks affecting the UK economy;
- ensuring greater consistency between fan charts for different variables;
- enabling us to produce a meaningful PSND fan chart for the first time;
- allowing us to pick out specific outcomes for further examination;
- and letting us assess a wider variety of fiscal rules, both individually and jointly.

1.3 As the first step in our stochastic simulation approach, we estimate a simple vector autoregression (VAR) model including the key drivers of PSND (such as the budget deficit, economic activity and interest rates), to which we append the debt accumulation identity that relates those drivers to PSND as a share of GDP. This structure can then be used to trace through the consequences of a shock to these variables, so producing a simulation. We draw our shocks from the vectors of historical residuals in the VAR (the estimated equations of the VAR do not fit exactly, leaving an unexplained 'residual'). Consequently, our simulations match the historical distribution of the residuals, including any skewness and 'fat tails', as well as replicating the empirical correlations between them. We repeat the process of drawing shocks and feeding them through the model thousands of times, with the thousands of simulations produced providing probability distributions for each variable that are consistent with the historical distribution of shocks affecting the UK economy.

1.4 In Box 4.2 in our October 2021 *EFO* we provided some initial results using stochastic simulations to assess the uncertainty surrounding the Government's proposed new fiscal targets, the derivation of which is described more fully in this working paper. Those results suggest that, on existing policies, there is a 54 per cent chance that PSND (excluding the Bank of England) falls as a share of GDP in 2024-25 (three years ahead from the date of that forecast) and a 61 per cent chance that the current budget is in surplus in 2024-25. These percentages are similar to those generated based on historical forecast errors, but yield wider distributions that likely better reflect the uncertainty around our central forecasts. We believe the stochastic simulation approach set out in this paper is now sufficiently developed for us to employ it in future *EFO*s as our primary method for calibrating fan charts and assessing fiscal rules. We nevertheless expect to continue to refine the approach in the future.

1.5 This paper is structured as follows:

- Chapter 2 describes how we currently portray the uncertainty around our central forecasts;
- Chapter 3 describes the general approach of stochastic debt simulations, how they are carried out by other organisations, and the specific details of our approach; and
- Chapter 4 describes the data used, the empirical results, and the use of stochastic simulations to assess the Government's fiscal rules.

2 How we currently depict uncertainty

2.1 In every *Economic and fiscal outlook* (*EFO*) we stress the uncertainty around our central forecasts for the economy and public finances. We do this in several ways, including:

- looking at past uncertainty, via **historical comparisons**;
- using probabilistic **fan charts based on historical forecast errors** for key macroeconomic and fiscal aggregates;
- making **comparisons between our forecasts and those of other external forecasters**;
- undertaking **sensitivity analysis**, to examine the fiscal implications of shocks to individual forecast determinants; and
- exploring the fiscal implications of a plausible combination of shocks to multiple forecast determinants via **scenarios**.

Historical comparisons

2.2 One of the simplest ways to illustrate the potential uncertainty surrounding the fiscal outlook is to examine the variation in long-run time series for key economic and fiscal aggregates. As the Bank of England's *A Millennium of economic data* dataset contains data on public debt and GDP going back over three centuries, we are able to capture a broad range of extreme events such as wars, previous pandemics, and severe economic depressions using this approach.

2.3 But historical comparisons provide a relatively crude tool for evaluating uncertainty. The structure of the economy and policy framework have changed dramatically during that period, so the response to a similar event that occurred in the distant past may not prove a good guide to what would happen today. And drawing inferences from the path of borrowing following specific shocks – like the aftermath of the 1918 flu pandemic – has its limitations, due to the different duration, economic impact, and implications for the public finances of the shock and also to other confounding factors (such as the aftermath of the first world war). To assess the uncertainty surrounding our forecasts, additional tools are therefore needed.

Fan charts based on historical errors

2.4 Another approach to capturing the uncertainty surrounding our forecasts involves using fan charts. The central forecast we produce is a median forecast, meaning that we expect an equal chance of the outturn eventually being above or below it. Our fan charts show four bands above and four below this central forecast, each of which represents a 10 per cent probability. So, if they accurately reflected the underlying uncertainty surrounding each central forecast at the time each central forecast was produced, then the outturn would fall within the bands of the fan 80 per cent of the time (and outside them 20 per cent of the time). And as the coronavirus pandemic has illustrated clearly, it is possible for very large shocks to occur from time to time and generate outcomes that lie well outside of our 80 per cent fans (Chart 2.1).

Chart 2.1: Fan chart for real GDP growth based on historical forecast errors

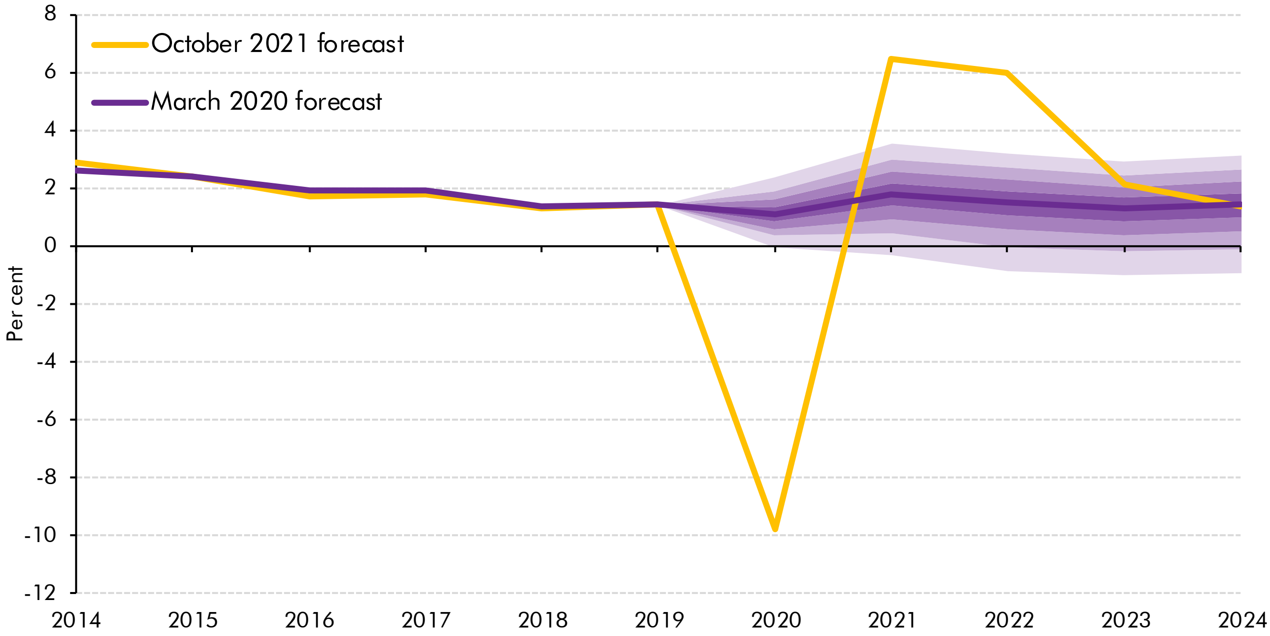

Source: ONS, OBR

2.5 To date, our fan charts have been calibrated using historical forecast errors, an approach which assumes future forecasts will be as accurate as past ones. Evidence from historical outturns suggests that adverse shocks (to both GDP and PSNB) are on average larger and therefore lead to larger forecast errors than favourable ones (technically the distribution is 'skewed' in an adverse direction). To capture this property, we generate the fans using a 'two-piece normal' distribution.[1] This distribution effectively splices together halves of two separate normal distributions, both with the same mode, but with different standard deviations.[2] These distributions are defined by three parameters: the mode and standard deviation and a skew parameter. The median is then matched to our central forecast. Together, the standard deviation and a measure of skew define the shape of the distribution, by determining the balance of risks, or the extent and degree to which risks are weighted to the upside or downside.[3]

[1] See Johnson, N., Kotz, S., and Balakrishnan, N., *Continuous univariate distributions, Vol. 1.*, 1994 and OBR, *Briefing paper No. 4 – How we present uncertainty*, June 2012, for further details.
[2] This distribution is commonly used by other forecasting institutions, including the Bank of England (for inflation). See, Britton, E., Fisher, P., and Whitley, J., *The Inflation Report projections: understanding the fan chart*, Bank of England, 1998.
[3] As shown in Banerjee, N., and Abhiman, D., *Fan chart: methodology and its application to inflation forecasting in India*, 2011, this skew indicator can be expressed as the difference between the mean and the mode of this distribution.

2.6 As discussed in a previous briefing paper, the errors used to generate the distributions for the real GDP historical-error-based fan charts are derived from official real GDP growth forecasts dating back to 1988.[4] Over much of this period errors have tended to be relatively closely clustered around zero. But forecast errors where the outturn is lower than forecast tend to be larger in magnitude than where the outturn is higher than forecast (as discussed in Annex A of our 2021 *Forecast evaluation report*, the large errors related to the pandemic have recently increased the mean absolute forecast error significantly). This reflects the distribution of growth itself, with the negative deviation from average growth rates during downturns greater than the positive deviation during upswings. As recessions are by their nature difficult to forecast, this negatively skewed distribution of the actual data carries over into an equivalent skewness in the distribution of forecast errors and therefore in the shape of the resulting fan, as shown in Chart 2.1.

2.7 We have also regularly used forecast errors for fiscal variables to create fan charts that illustrate the uncertainty surrounding our fiscal forecasts and enable us to assess the probability of the Government meeting its fiscal rules. For instance, for PSNB as a share of GDP, we have been able to use errors relating to forecasts published from 1988 onwards (although for some other fiscal variables only somewhat shorter time series are available).

2.8 However, even though it is of great interest to policymakers, to date we have not produced a fan chart for the level of debt-to-GDP. In principle, we could adopt the same historical forecast error approach that we use for PSNB fan charts, but in practice this would require a very long run of history to capture adequately the observed range of shocks to debt-to-GDP. An alternative approach of deriving the distribution of debt shocks indirectly from the distribution of historical forecast errors for annual borrowing – which drive the accumulation of debt – would suffer from the problem that those errors are serially correlated i.e. higher than expected borrowing in one year is often followed by similar errors in the next few years. But it is feasible to derive a debt-to-GDP fan chart indirectly from what we know about the uncertainty surrounding the deficit, by capturing the persistence of shocks to debt directly, via the stochastic simulations approach explained in this paper.

2.9 A fan chart surrounding our March 2020 PSNB-to-GDP forecast, produced using the historical forecast errors method described above, is shown in Chart 2.2.[5] It shows that the impact of the pandemic pushed PSNB far outside the range implied as 80 per cent probable by the fan. As adverse shocks like this raise PSNB, the fan has an upward skew (mirroring the downwardly skewed GDP growth errors shown in the fan above – i.e. both are skewed in an adverse direction). In addition, it shows the spread of PSNB errors widening rapidly as the forecast horizon increases past two years. This widening reflects the fact that it is the level of GDP (as opposed to the growth rates shown above) that is the key determinant of the level of PSNB and that, in the absence of policy changes, forecast errors in growth rates are likely to be serially correlated. Countering this, one might expect policymakers to act if

[4] OBR, *Briefing paper No. 4 – How we present uncertainty*, June 2012.

[5] Due to difficulties associated with identifying the modal error for PSNB as a share of GDP, a slightly more complicated approach to producing fan charts was set out in Briefing paper No. 4, and used in some of our *Economic and fiscal outlooks* – this involved using ready reckoners to infer a modal forecast error for PSNB from the modal error for GDP growth. We have also adjusted fiscal fan charts for other reasons, such as to reflect a more symmetrical degree of uncertainty surrounding cyclically-adjusted fiscal aggregates, or by imposing judgement to ensure the fans widen sufficiently at the forecast horizon to reflect the prevailing degree of uncertainty.

borrowing rises sufficiently – and this is likely to prevent the fan from growing too wide. As policy takes time to be implemented this gives the fan a distinctive 'wine bottle' shape over the five-year forecast period.

Chart 2.2: Fan chart for PSNB based on historical forecast errors

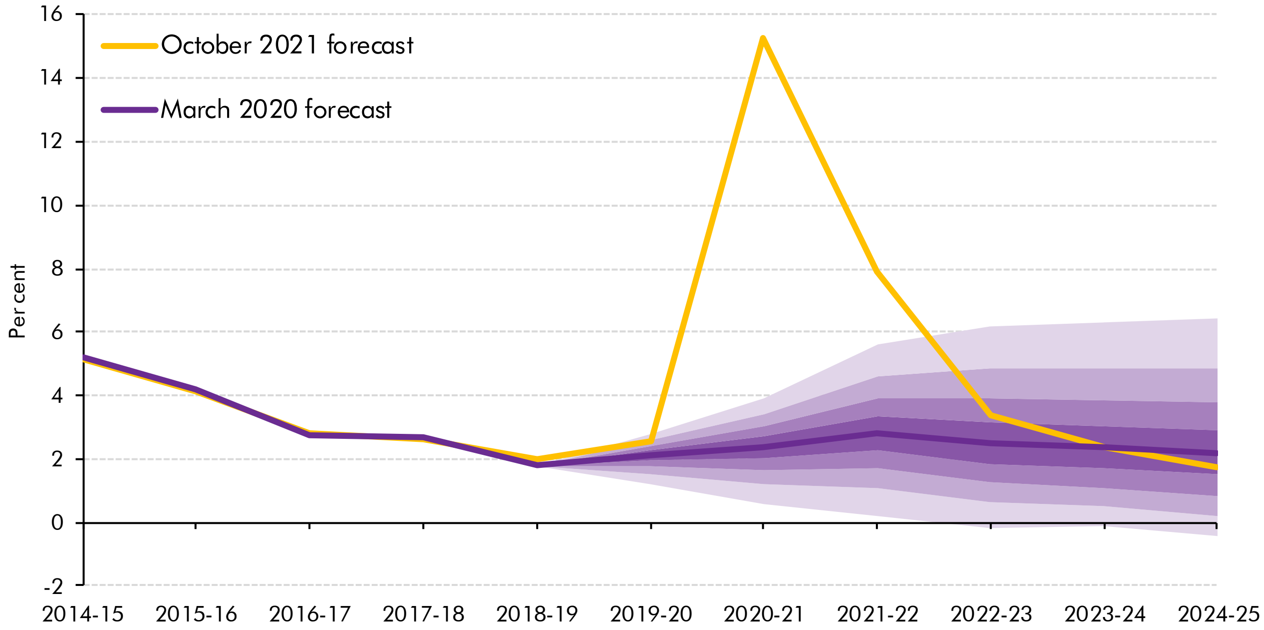

Source: ONS, OBR

Comparing our forecast with other forecasters

2.10 We regularly include a comparison against external forecasts in our *EFOs*. Looking at the range of forecasts is not a guide to uncertainty around our forecast, as it measures the dispersion of central opinions, while each forecaster will in fact have their own expected distribution around their stated central forecast. So instead these forecast comparisons generally provide a guide to the disagreement between forecasters, usually on a few key judgements made in the central forecast (such as the underlying trend in potential output). It seems plausible, however, that there is more room for disagreement when uncertainty is high than when uncertainty is low. This was evident in the early months of the pandemic with the spread of GDP forecasts for 2020 widening considerably.

Sensitivity analysis

2.11 In Chapter 4 of our *EFOs*, we illustrate the effect of individual sources of uncertainty surrounding our forecasts via sensitivity analysis. This approach involves assessing by how much key economic and fiscal judgements would have to vary for the Government to fail its fiscal rules. In our October 2021 *EFO*, we quantified the relatively small changes in potential output, effective interest rates, RPI inflation, and current departmental spending required to wipe out the relatively modest headroom the Government had reserved against its current budget balance target in 2024-25. We also looked at the changes in primary borrowing, nominal GDP growth and interest rates, and financial transactions that would prevent it hitting its debt-falling target in the same year. But while this approach usefully illustrates the impact of alternative assumptions, because it does not provide a corresponding distribution for those assumptions, it cannot provide an overall measure of the uncertainty surrounding our forecasts.

Scenarios

2.12 In our *EFOs*, we normally set out the fiscal implications of some illustrative alternative economic scenarios. These scenarios are also designed to highlight the sensitivity of our central forecast to some of the key judgements underpinning it. To produce the scenarios, we sometimes make use of alternative models, such as our small macroeconomic model, together with stylised assumptions as to how a change in a key assumption (such as trend productivity growth) feeds through into other parts of the economy, and then affects the fiscal forecast.[6] Recent scenarios have included upside and downside coronavirus scenarios in our November 2020 *EFO* (Chart 2.3), and higher and more persistent inflation scenarios in our October 2021 *EFO*. But this approach is subject to the same drawbacks as sensitivity analysis – there is not usually a straightforward way to ascribe probabilities to the scenarios, so interpreting their implications for the uncertainty surrounding the forecast is far from straightforward. In addition, scenarios are relatively resource-intensive to produce, so it is not possible to produce ranges of scenarios that fully encapsulate all the ways outcomes could differ from our central forecast in sufficient detail to be worthwhile at the same time as completing a Budget forecast.

Chart 2.3: Coronavirus scenarios for real GDP from our November 2020 *EFO*

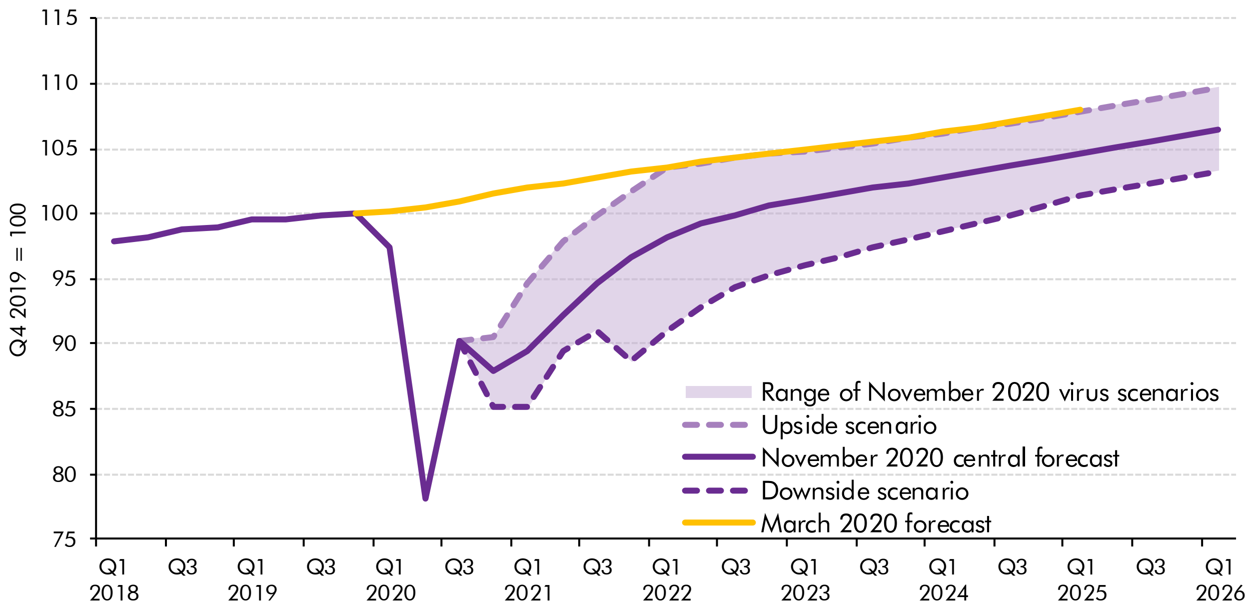

Source: ONS, OBR

[6] Our small macroeconomic model is an updated version of the model presented in Murray, J., *Working paper No. 4 – A small model of the UK economy*, OBR, July 2012.

3 Stochastic simulations

3.1 As the next step in the portrayal of uncertainty around our central forecast, we plan to employ the method of 'stochastic simulation'. [7] This is already widely used in the academic literature and by some other official organisations, such as the IMF. This chapter: sets out the general modelling choices for stochastic simulations; details the specific choices involved when simulating public sector net debt (PSND) as a share of GDP; reviews the use of stochastic debt simulations by other organisations and in the academic literature; and describes our implementation of the method.

What are stochastic simulations?

3.2 Stochastic simulations are an extension of the scenario method used to generate a probability distribution for a variable (or set of variables) of interest. This is accomplished by first constructing or assuming a model that describes the interrelationships between the variables of interest (for instance, output, inflation, interest rates and the budget deficit). Typically, these relationships will not fit exactly – there will be unexplained 'residuals' or shocks. Alternative scenarios can then be generated by assuming different plausible values for these shocks and tracing their consequences through the model. Assuming that the set of scenarios that are constructed is typical of those likely to be encountered (i.e. that the distribution of hypothetical shocks matches the actual distribution), by generating a very large number (several thousand) of such hypothetical scenarios, the corresponding probability distributions for the variables of interest can then be computed. The probability of a particular event (or set of events) happening can also be calculated merely by observing the fraction of these hypothetical scenarios in which it occurs.

3.3 There are several decisions to make when carrying out stochastic simulations with multiple variables:

- How to **capture correlations between shocks to variables, and correlations over time.** One approach (that we also adopt) is to use a multivariate model such as a vector autoregression (VAR) model to describe the dynamic relationship between the variables of interest. In a VAR, each variable is driven by an estimated equation that contains lags of itself and lags of all the other variables as explanatory variables, together with an unexplained residual (here also the 'shock'). These residuals will be uncorrelated over time by construction, though the residuals in different equations may be contemporaneously correlated. Scenarios (or simulations) can then constructed by drawing vectors of shocks from a multivariate probability distribution with the same statistical characteristics (e.g. variance-covariance matrix) as the historical residuals and then using the model to trace their effects over time on the variables of interest.

[7] As set out in Box 4.2 of our October 2021 *Economic and fiscal outlook*.

This approach therefore replicates any correlations between the shocks. Sometimes – for instance when there are insufficient data available to estimate a VAR – it may be necessary to shock the variables of interest directly, using the empirical distribution of the raw data rather than the estimated residuals, though in this case additional adjustments would be necessary to allow for any serial correlation in the shocks (this approach is tantamount to employing a VAR that contains only constants and no explanatory variables).

- **How to randomly generate the shocks**. Whether using a VAR or applying shocks directly to the raw data, one possibility is to assume the shocks conform to a particular distribution, such as a joint normal distribution. But their empirical counterparts (the residuals in the case of the VAR) may not in fact be normally distributed. Indeed, many economic variables exhibit both 'fat tails' (extreme outliers seem to occur more often than they would under a normal distribution) and are skewed ('bad' outcomes are more likely than 'good' outcomes). So, an alternative approach that replicates these features is to draw the shocks by sampling repeatedly and randomly from the set of estimated historical residuals in the VAR (known as 'bootstrapping'). This has the advantage that the resulting simulations will necessarily replicate the historical distribution of the residuals, including any skewness and 'fat tails'.

- How to **generate the central forecast**. If the same model is used to generate the central forecast as to evaluate uncertainty, then this will drop out directly from the stochastic simulations. For instance, the mean, median and modal outcomes of the set of stochastic simulations can all be easily obtained. For many uses, especially in academic research, this is quite sufficient. However, at the OBR (and at other official forecasting bodies) the central forecast is based on a more complicated macroeconomic model than a VAR and incorporates much 'expert' judgement that draws on multiple sources of additional information. We will continue to use this approach to construct our central forecasts, while using stochastic simulations to guide our assessment of the uncertainty surrounding those central forecasts. Consequently, we need to incorporate an additional step that shifts the distributions obtained from the stochastic simulations so that the associated median outcomes coincide with our in-house central forecast.

- Whether to allow for **uncertainty around coefficient estimates** (a form of 'model uncertainty'). The coefficients in the model used to generate the shocks and map them into the variables of interest are generally not known with certainty and must instead be estimated. If something is known about the precision with which such coefficients are estimated (e.g. from the standard errors of the estimated coefficients) then allowance can be made for this additional source of uncertainty by also randomly drawing alternative sets of coefficients. (The standard approach outlined above can be thought of as randomly shocking just the constants in the VAR in each period; shocking the other coefficients too can therefore be seen as a natural extension, though it complicates implementation somewhat.)

Using the stochastic simulation approach for PSND

3.4 As discussed in the previous chapter, one advantage of stochastic simulations is that they allow us to produce a probability distribution around our forecasts for PSND. The standard stochastic simulation approach to forecasting PSND also works by generating a very large number of different PSND simulations, but these are not generated by directly applying stochastic shocks to an underlying PSND forecast. Instead thousands of scenarios are generated for key drivers of debt-to-GDP (such as interest rates, GDP growth and inflation), typically generated by using a VAR to capture the interdependencies between them. The simulations for PSND are then generated by plugging the scenario debt driver into a debt accumulation identity like that shown below (equation (1)). The equation describes how PSND as a share of GDP evolves from one period to the next, here based on the effective interest rate on that debt (EIR), real GDP growth (g), inflation (π), the primary deficit (PD), and 'stock-flow adjustments' (SFAs). SFAs are changes in the stock of debt that are not accounted for by flows of primary borrowing and debt interest; they can, for example, arise from the acquisition of financial assets. In each simulation, paths for the determinants of debt – real GDP growth, interest rates, inflation, the primary deficit and SFAs – are generated and equation (1) is then used to calculate the associated path for PSND.

$$\frac{PSND_t}{GDP_t} = \frac{(1 + EIR_t)}{(1 + \pi_t)(1 + g_t)} x \frac{PSND_{t-1}}{GDP_{t-1}} + \frac{PD_t}{GDP_t} + \frac{SFA_t}{GDP_t} \qquad (1)$$

3.5 The primary deficit can be embedded within the VAR along with all the other variables of interest. The equations of a VAR are, however, typically not structural in nature. Instead, they are 'reduced forms' that are convolutions of the underlying structural equations of the model economy. Extracting these underlying structural equations requires the imposition of additional 'identifying' assumptions (there is a large academic literature on how to do this). If, as is sometimes the case, researchers want to be able to investigate the consequences of following alternative fiscal rules, then the VAR equation for the primary deficit needs to be replaced by a structural 'fiscal reaction function' that can be altered appropriately. The difference is that the latter may include contemporaneous variables, exclude some lagged variables, etc. It may also include past values of debt. As the OBR is forbidden from analysing alternative policies, our present approach simply embeds the primary deficit within the VAR.

3.6 One driver of debt that is not (typically) included in the VAR is the SFA term. So instead to capture its effect on debt another equation or assumption can be used to generate scenarios for it, with the values then plugged into the debt-accumulation equation. However, the SFA is often ignored altogether in stochastic debt simulations for a couple of reasons. Firstly, sizable SFA movements are often driven by hard to predict one-off events which make them difficult to model. And secondly, over extended periods SFAs (as a share of GDP) typically average out close to zero so do not significantly affect the median debt projection. However, this ignores that they affect the width and shape of the future distribution of debt - while the IMF have shown that large spikes in public debt are often driven by sizable SFAs highlighting the importance of trying to capture them.[8] In the UK, while the average SFA (as a share of GDP) has been relatively close to zero, they are not normally distributed so would affect the width and skew of the fan (Chart 4.2). There have been several very large positive SFAs (pushing up debt) over the past several decades, typically during periods of stress on the UK economy. This is because governments frequently step in to take over systemically (or electorally) important distressed assets during bad times (such as the injections of capital into banks during the financial crisis).

3.7 There are several advantages to using stochastic simulations compared to our current historical error fan chart approach:

- A limitation of our current approach to constructing fan charts is that the forecast errors have only a limited back series and so may constitute an unrepresentative sample. In particular, they may not include a sufficient number of large but rare shocks ('catastrophic risks'). Using stochastic simulations, we can capture a longer and hopefully more representative back history covering several decades or more (our preferred specification starts in the mid-1950s). We can also more easily calibrate fans against particular sub-periods of interest, such as the unusually benign period running from 1993 to 2006 (the 'Great Moderation') or the unusually volatile past 13 years containing two once-in-a-century shocks.

- The stochastic simulation approach ensures consistency between the fan charts for different variables, particularly borrowing and debt. And by drawing vectors of shocks from the corresponding historical distribution it better captures the correlations present in the data than our current historical error approach.

- The stochastic simulation approach also allows us to produce a meaningful distribution for the level of PSND for the first time, as it captures the way borrowing feeds through to the stock of debt.

- As with our scenarios, the approach also allows one to pick out specific realisations for further examination and discussion.

- Finally, it allows us to calculate the probability of meeting different targets more easily, as well as facilitating the calculation of the probability of meeting multiple targets simultaneously.

3.8 The main drawback of the stochastic simulation approach relative to our existing approaches to portraying uncertainty is that it is more complicated and requires greater technical expertise to apply. It is perhaps also harder for non-expert users to grasp. For this reason, we will continue to use a variety of approaches to illustrate uncertainty around our central forecasts but plan to use the stochastic simulation approach as our primary method for generating fan charts and assessing fiscal rules.

[8] L. Jaramillo, C. Mulas-Granados and E. Kimani, *The blind side of public debt spikes*, IMF working paper October 2016.

How others use stochastic debt simulations

3.9 Stochastic debt simulations are already used by several other official institutions to create fan charts. In this section, we provide an overview of the methods used by international institutions and other independent fiscal institutions (IFIs), as well as in the academic literature.

3.10 The IMF uses stochastic simulations in its debt sustainability analysis for what it terms 'market access countries', such as the UK. To date, it has calculated the variance-covariance matrix using the raw data for the drivers of the debt-to-GDP ratio and assumes shocks have a joint normal distribution.[9] The IMF has recently published an updated methodology, in order to move away from the assumption that shocks are normally distributed and to better capture asymmetric risks around the baseline.[10] It is a two-step process to produce 'realism-adjusted' fan charts around the its baseline debt-to-GDP forecast. First, the IMF generates a 'historical' fan chart from the joint distribution of the historic raw data for the key debt-to-GDP drivers. To capture the persistence of shocks, the stochastic realisations are drawn from this distribution using a 'block-bootstrap' approach (each shock is drawn for a consecutive two-year block). Second, the 'historical' fan chart is compared to the IMF's baseline forecast and adjusted if necessary.[11]

3.11 The European Commission (EC) also uses stochastic simulations as part of its debt sustainability analysis.[12] For each country, it calculates the historical variance-covariance matrix using the raw data for five key drivers of debt and assumes a joint normal distribution to generate the shocks. It constructs fan charts by applying these shocks to its baseline forecast for each variable to generate different paths for the debt-to-GDP ratio.

3.12 As part of its debt sustainability analysis, the European Central Bank (ECB) also uses stochastic debt simulations.[13] Its approach differs from the IMF and EC's as it first estimates a VAR. The historical shocks are the residuals from this VAR and it obtains the distribution of shocks using bootstrapping (with replacement). The ECB then forecasts each variable using the estimated VAR coefficients and shocks randomly drawn from the historical distribution. Finally, it combines these forecasts with a fiscal forecast (which depends on the real GDP forecast) to generate a distribution of debt-to-GDP paths for a fan chart.

3.13 In addition to international institutions, the OECD's review of the OBR noted that stochastic debt simulations are already used by several other IFIs.[14] Similarly, a recent survey found

[9] IMF, *Modernizing the framework for fiscal policy and public debt sustainability analysis*, August 2011.
[10] IMF, *Review of the debt sustainability framework for market access countries*, January 2021.
[11] If the baseline forecast for the debt-to-GDP ratio is inside the 20th percentile of the historical fan chart, it creates a 'standard' (symmetric) fan chart by applying the demeaned shocks to the baseline forecast. If the baseline forecast is outside the 20th percentile of the historic fan, it instead creates a 'realism-adjusted' (asymmetric) fan chart by adjusting the centre of the fan chart using historical cross-country comparisons.
[12] European Commission, *Debt sustainability monitor 2020*, February 2021. The method is based on Berti, K., *Stochastic public debt projections using the historical variance-covariance matrix approach for EU countries*, April 2013.
[13] Bouabdallah, O., et al, *Debt sustainability analysis for euro area sovereigns: a methodological framework*, ECB Occasional Paper Series, October 2017.
[14] OECD, *OECD independent fiscal institutions review: Office for Budget Responsibility of the United Kingdom*, September 2020.

that 8 out of 15 EU IFIs carry out debt sustainability analysis using stochastic simulations.[15] Some use the methodologies discussed above, for example Slovenia's two IFIs and the Lithuanian National Audit Office use the current IMF method and the Italian Parliamentary Budget Office has used the similar EC method to report stochastic debt-to-GDP projections.[16] Other IFIs use proprietary methods. The OECD's report highlighted the Spanish Independent Authority for Fiscal Responsibility's work using a VAR-based model and bootstrapping to assess the likelihood of meeting Spain's debt target.[17] Finally, the Irish Fiscal Advisory Council's new small scale structural model was used to create fan charts for assessing fiscal rules in its 2021 Fiscal Assessment Report.[18] The model has three key equations covering growth, the Phillips curve and the marginal interest rate and is solved stochastically to create fan charts for the Council's debt projections.[19]

3.14 Stochastic debt simulations have also been widely used in the academic literature. Table 3.1 provides a summary of approaches from a selection of papers. The majority of which generate shocks to (economic) determinants using VARs. The papers include similar (economic) variables within the VAR to us, namely real GDP growth, interest rates and inflation. Some papers include both short and long-term interest rates or include foreign interest rates and exchange rates (important if the country has significant borrowing in foreign currency). Many papers incorporate a structural fiscal reaction function for the primary deficit instead of embedding it within the VAR. Most papers also ignore SFAs completely, assuming they average out to zero.

[15] The Network of EU Independent Fiscal Institutions, *The role of the Independent Fiscal Institutions in assessing the sustainability of high public debt in the post-Covid era*, February 2021.
[16] Ufficio Parlamentare di Bilancio, *2020 Budgetary Policy Report*, December 2019.
[17] For example see AIReF, *Monthly stability target monitoring 2021*, July 2021. More detail on methodology can be found in Cuerpo, C., *Spanish public debt sustainability analysis*, AIReF Working Paper, October 2014.
[18] Irish Fiscal Council, *Fiscal Assessment Report*, May 2021.
[19] Casey, E. and D. Purdue, *Maq: A Fiscal Stress Testing Model for Ireland*, Irish Fiscal Advisory Council Working Paper Series No. 13, February 2021.

Table 3.1: Summary table of a selection of papers in the academic literature

	Paper	Data	How are shocks generated?	FRF used?	SFA included?
1	Stochastic public debt projections using the historical variance-covariance matrix approach for EU countries (2013) by Berti	Economy data quarterly, fiscal annual for 24 EU countries	Variance-covariance matrix of historical shocks, assuming joint normal distribution	Estimated coefficients of sensitivity of primary balance to economic cycle used	No
2	Debt sustainability analysis for euro area sovereigns: a methodological framework (2017) by Bouabdallah et al	Economy data quarterly, fiscal annual for Euro Area	VAR, bootstrapping	Yes	No
3	Assessing Dutch fiscal and debt sustainability (2020) by Carton and Fouejieu	Economy data quarterly, fiscal annual for Netherlands	VAR, bootstrapping	Yes	No
4	Public Debt Dynamics: The Effects of Austerity, Inflation, and Growth Shocks (2012) by Cherif and Hasanov	Quarterly data for US	VAR, bootstrapping	Primary balance included in VAR	No
5	Spanish Public Debt Sustainability Analysis (2015) by Cuerpo and Ramos	Quarterly data for Spain	VAR, bootstrapping	Components of primary balance in VAR	No
6	Stochastic debt simulation using VAR models and a panel fiscal reaction function: results for a selected number of countries (2012) by Medeiros	Quarterly data for 15 EU countries	VAR, both bootstrapping and joint normal distribution	Yes	No
7	When is debt sustainable? (2012) by Lukkezen and Rojas-Romagosa	Annual data for US and 6 EU countries	VAR, joint normal distribution	Yes	No
8	Stochastic forecast of the Slovak public debt (2016) by Výškrabka	Economy data quarterly, fiscal annual for Slovakia	VAR, bootstrapping	Yes	Yes modelled as iid process

OBR approach to stochastic simulations

3.15 Our first step in producing stochastic simulations for PSND is to estimate a VAR for the key drivers of deficits and debt. Our VAR includes quarterly data for real GDP growth, inflation, interest rates, and the primary deficit (instead of a separate structural fiscal reaction function). The VAR therefore incorporates the average historical response of both fiscal and monetary policies to economic developments, as well as the impact of those policies on the economy.

3.16 A feature of our VAR plus debt-accumulation equation model worth noting is that we include two interest rates, rather than one, in order to capture changes in the effective maturity of the debt of the consolidated public sector (government plus Bank of England). In recent years, the effective maturity of UK public debt has shortened as a consequence of the asset purchases by the Bank of England's Asset Purchase Facility (APF). [20] As of the end of October 2021 some £740 billion of gilts are held by the APF, paid for by the issuance of bank reserves on which the Bank pays Bank Rate. [21] The surpluses/deficits on the APF also ultimately accrue to the Treasury. An increase in Bank Rate thus straightaway raises public sector interest costs, instead of feeding through slowly over several years. This means that the effective interest rate (EIR) on PSND will be more responsive to changes in Bank Rate than it was in the past. We capture this by including both Bank Rate (BR), to capture APF interest, and an EIR that excludes the effects of APF (EIRx) in the VAR and then later add together APF and non-APF interest in our debt accumulation equation.

3.17 The second step is to draw vectors of shocks from the VAR. We draw our shocks randomly (with replacement) from the vectors of historical estimated residuals in the VAR. Consequently, our simulations match the historical distribution of the residuals, as well as replicating the empirical correlations between them. For each period of each simulation, we draw the shocks from a suitable sample period; our preferred period extends from the mid-1950s (when most quarterly data are first available) to the present. But we can vary the sample period from which shocks are chosen in order to illustrate the consequences of different assumptions about the degree of prevailing uncertainty; for example, a sample period spanning just the 'Great Moderation' would produce much narrower fans. We also allow for coefficient uncertainty, by adding shocks to the coefficients of the VAR.[22] We then repeat this process to generate several thousand simulations (this paper uses 10,000), allowing us to generate probability distributions for each of the variables.

3.18 Our third and fourth steps are to generate simulation paths for the SFA term. We model this using a simple regression that relates it to a constant and a dummy variable for whether UK real GDP growth is negative in that period of the simulation. We then generate scenarios for SFAs using the coefficients of the regression, real GDP growth in each scenario as an input, and again drawing our shocks randomly (with replacement) from the estimated residuals.

3.19 In our final steps we bring together the simulated debt determinants from the VAR and the SFA and use the debt accumulation equation to calculate simulated paths for PSND. To make this calculation, we combine our central forecast for the size of the APF (a possible extension would be to endogenise this) with the simulated path for Bank Rate to calculate APF interest. We also calculate other fiscal variables of interest such as PSNB. In using the results to calibrate the uncertainty around our central forecast, we add an adjustment to align the median of the fans with our central forecast (this is described in paragraph 3.20). Finally, we construct the fans using the probability distributions that result from the multiple simulations. Figure 3.1 shows a diagram of the key steps in our approach.

[20] See Box 4.5 of our July 2021 *Fiscal risks report*.
[21] This is the redemption value of APF gilts, which we use in our model, rather than the purchase value which was close to £840 billion at the end of September 2021. Gilts are valued at their redemption value in the PSND calculation.
[22] We do this by drawing from a multivariate normal distribution using the covariance matrix of the estimated coefficients.

Figure 3.1: Diagram of approach

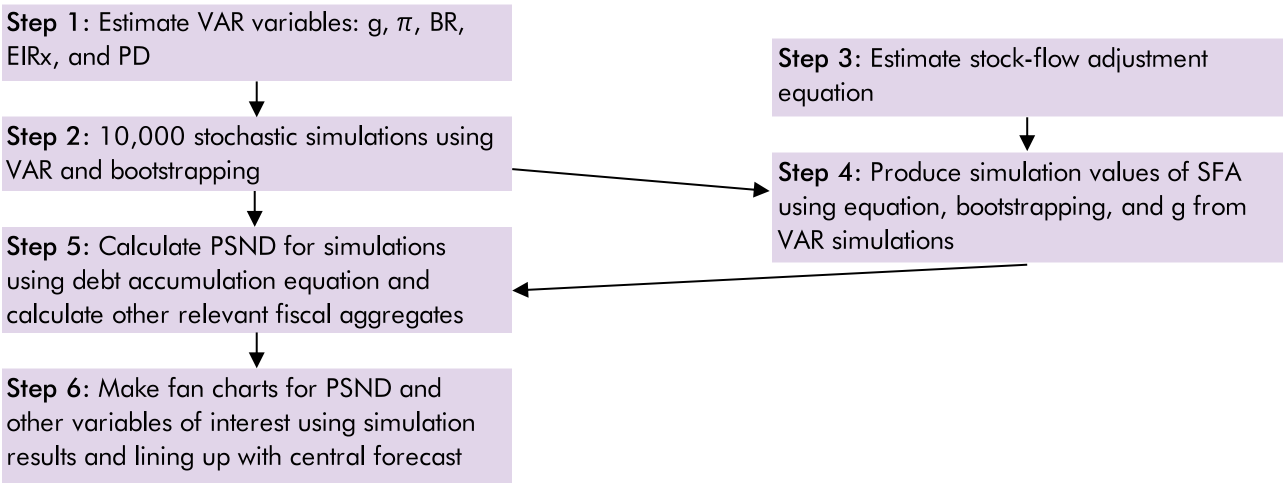

Lining up the fan charts with our central forecast

3.20 To use the stochastic simulations to assess the probabilities around our central forecast and of meeting fiscal rules in our *EFO* we need to align the median of the distributions with our central forecast. Our central forecast embodies our own forecast judgements and explicitly factors in current stated policies. Once this is done it means that the median/centre of the fan reflects our own judgements, while the size and shape of the bands around it are determined by the long run of history captured in the VAR. We align the median in two steps. Firstly, we include a fixed set of residuals in the VAR and the SFA equation over the forecast period that align the VARs deterministic solution to our central forecast. But as the VAR errors are skewed this will not exactly match the median stochastic solutions so we incorporate a further small manual adjustment at the end to align them exactly. (We need to include these fixed residuals in the VAR at the outset as there is a non-linearity in the model for the SFA.)

4 Data and results

4.1 Our VAR comprises quarterly data for real GDP growth, GDP deflator inflation, Bank Rate (to capture APF interest), the effective interest rate on PSND (excluding the effects of the APF), and the primary surplus as a share of GDP. The series start in 1955 (although for the EIR, we have calculated a back series by using the central government net cash requirement – the main determinant of debt – to create a quarterly series for PSND pre-1993), and all are seasonally adjusted. Time series plots are shown in Chart 4.1 and further details of all series used are provided in Table A.1 in the annex. Standard unit root tests on the variables suggest they are stationary apart from the interest rate variables where we are unable to reject a unit root (see annex Table A.2). For consistency with inflation, we have chosen to keep interest rates in levels in the VAR.

Chart 4.1: Data series used in the VAR

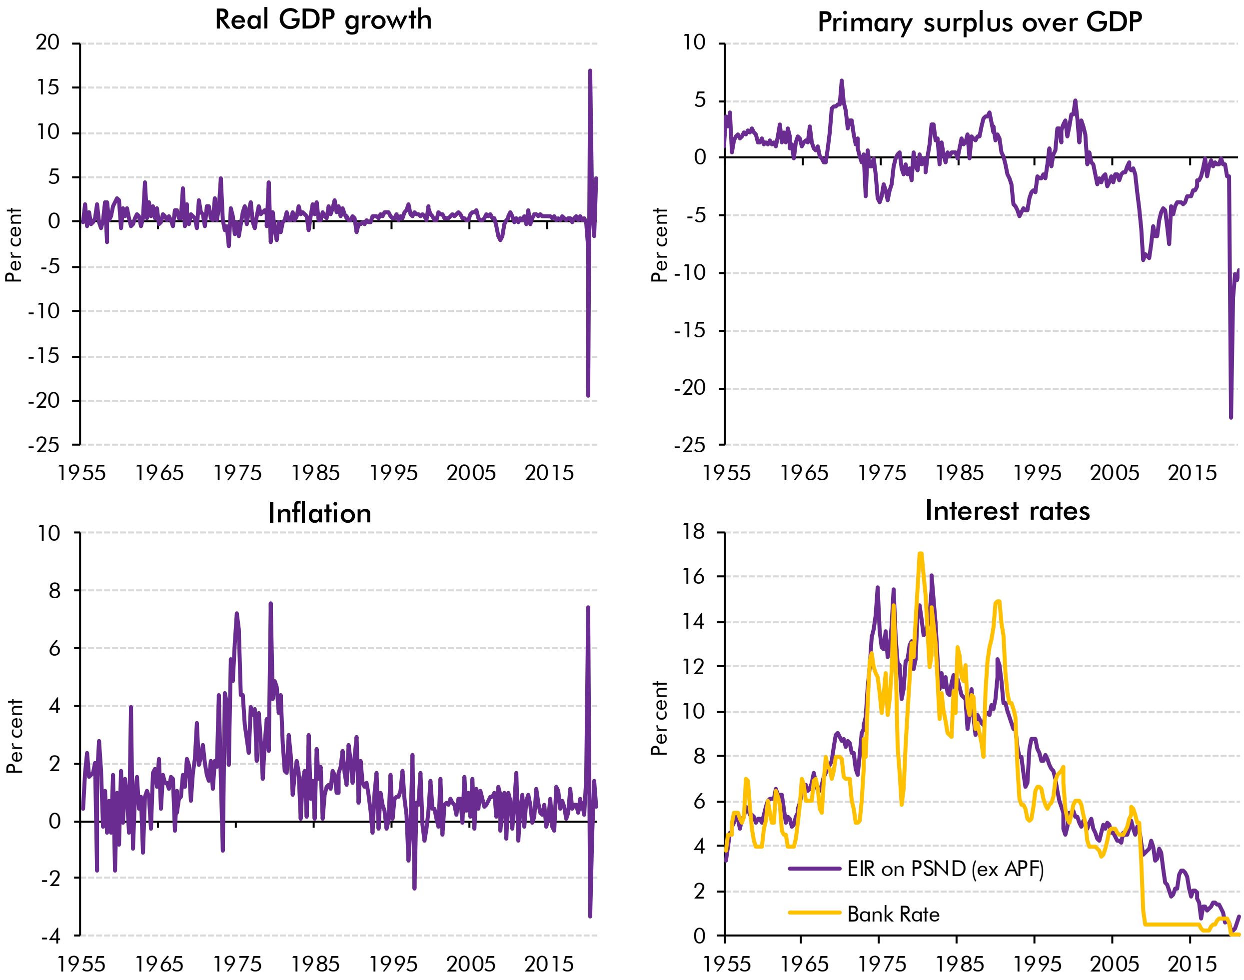

Source: ONS, OBR

4.2 We use financial year data for the stock-flow adjustment term, starting in 1955-56. Historically the average SFA as a share of GDP has been close to zero. Chart 4.2 shows the time series on the left-hand side and the corresponding distribution on the right-hand side. While the average is close to zero there are a few very large positive values in 1971-72 and 1974-75, during and after the financial crisis in 2008-09 and 2009-10, then again in 2016-17 and in 2020-21. There is thus clear evidence of non-normality, which is confirmed by formal tests using the Jarque-Bera statistic.[23]

Chart 4.2: Stock-flow adjustment time series and distribution

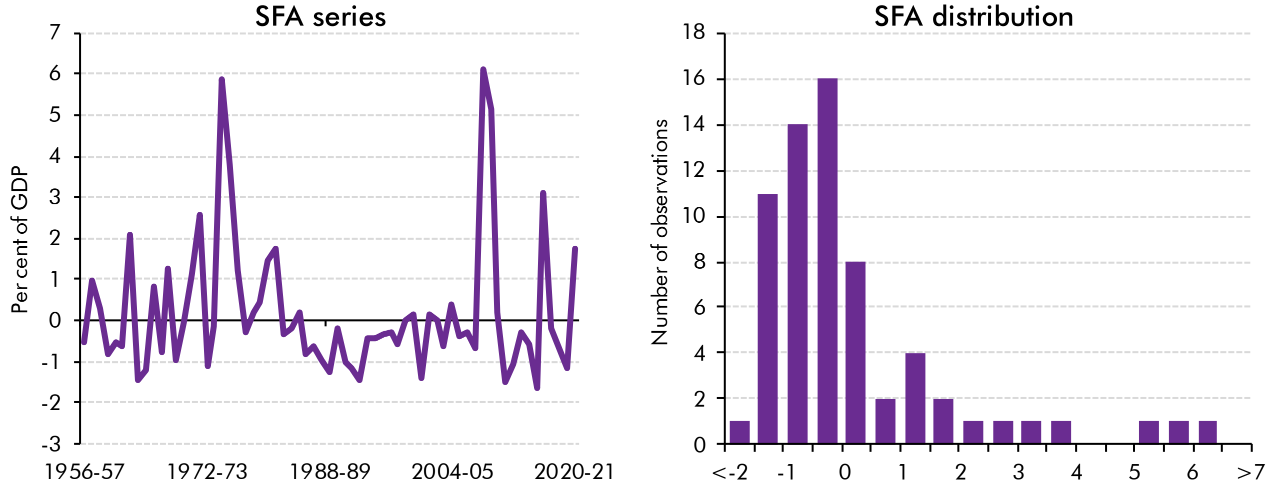

Source: ONS, OBR

Results

4.3 The sample used to estimate our VAR runs from the third quarter of 1956 to the final quarter of 2019. This means the VAR coefficients themselves are not distorted by the extreme movements in activity and the budget deficit during the pandemic. But when generating the stochastic simulations we *do* include the implied residuals during the pandemic period, i.e. up to the second quarter of 2021. The annex provides further details on the results and post-estimation tests. To remove serial correlation from the residuals we include five lags of the variables in the VAR (Table A.3). All the roots are within the unit circle (although one is close to unity, corresponding to the near unit-root behaviour of interest rates) (see Chart A.1). Normality of the residuals is rejected in formal tests, supporting the use of bootstrapping rather than drawing from a normal distribution in conducting the simulations (Table A.4). For the SFA equation, we found that the best fit came from a simple equation that included a constant term and a dummy variable which indicated whether or not real GDP was falling (Table A.5).

4.4 The panels of Chart 4.3 show fans for real GDP growth obtained by stochastic simulation and our usual historical forecast errors approach. The median (centre of the fan) for both approaches has been set to match our October 2021 forecast. Between 2023 and 2026 the width of the fan is similar in both approaches, with the fan based on historical errors having

[23] The Jarque-Bera statistic is 85.04, with a P-value of 0.000, implying that normality is rejected at the 1 per cent significance level.

a slightly more pronounced downward skew (bottom panel of Chart 4.3). But in the first two years, the stochastic simulation fan is wider, with 2022 being the widest year of the fan.

4.5 The inclusion of model uncertainty in our stochastic simulation approach widens the GDP fan, particularly in the first two years. Model uncertainty depends on the values of the explanatory variables (as the randomly drawn coefficients are multiplied by them), with large outliers in the explanatory variables generating greater model uncertainty. At the start of our simulation period (the third quarter of 2021) the lags in the VAR will include outturn data from over the pandemic period (this includes the second quarter of 2020 where GDP fell by almost 20 per cent). Excluding model uncertainty, would narrow the GDP fan in 2022 from 8 percentage points to a little over 5 percentage points (like our historical error approach). In contrast, in later years when median quarterly GDP growth is stable, coefficient uncertainty is smaller, adding only around 1 percentage point to the fan width.

Chart 4.3: GDP fan charts

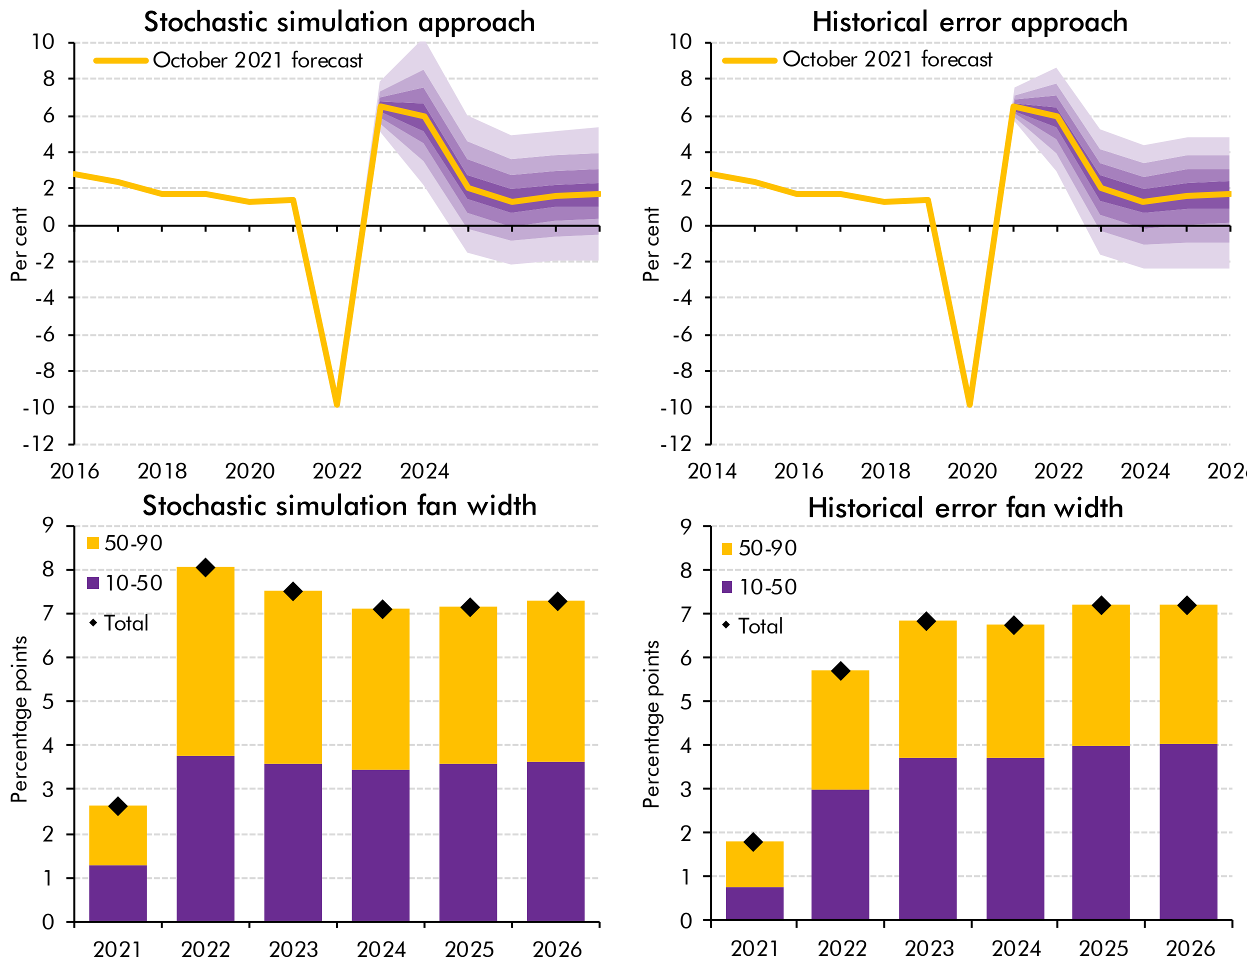

Source: ONS, OBR

4.6 Chart 4.4 shows public sector net borrowing (PSNB) fans using the stochastic simulation and historical error approaches, with the medians again aligned with our October 2021 central forecast. Both sets of fans have an upward skew, but the stochastic simulation fan is again somewhat wider in 2021-22 and by the final year 2026-27 although similar in the years in between. The effect of including model uncertainty in the stochastic simulations also increases the width of the fan. As with GDP growth, this effect is larger in the early years although the difference is less pronounced than for GDP growth (as the effect of model uncertainty on interest rates, another driver of PSNB and PSND, is similar in each year of the forecast).

Chart 4.4: PSNB fan charts

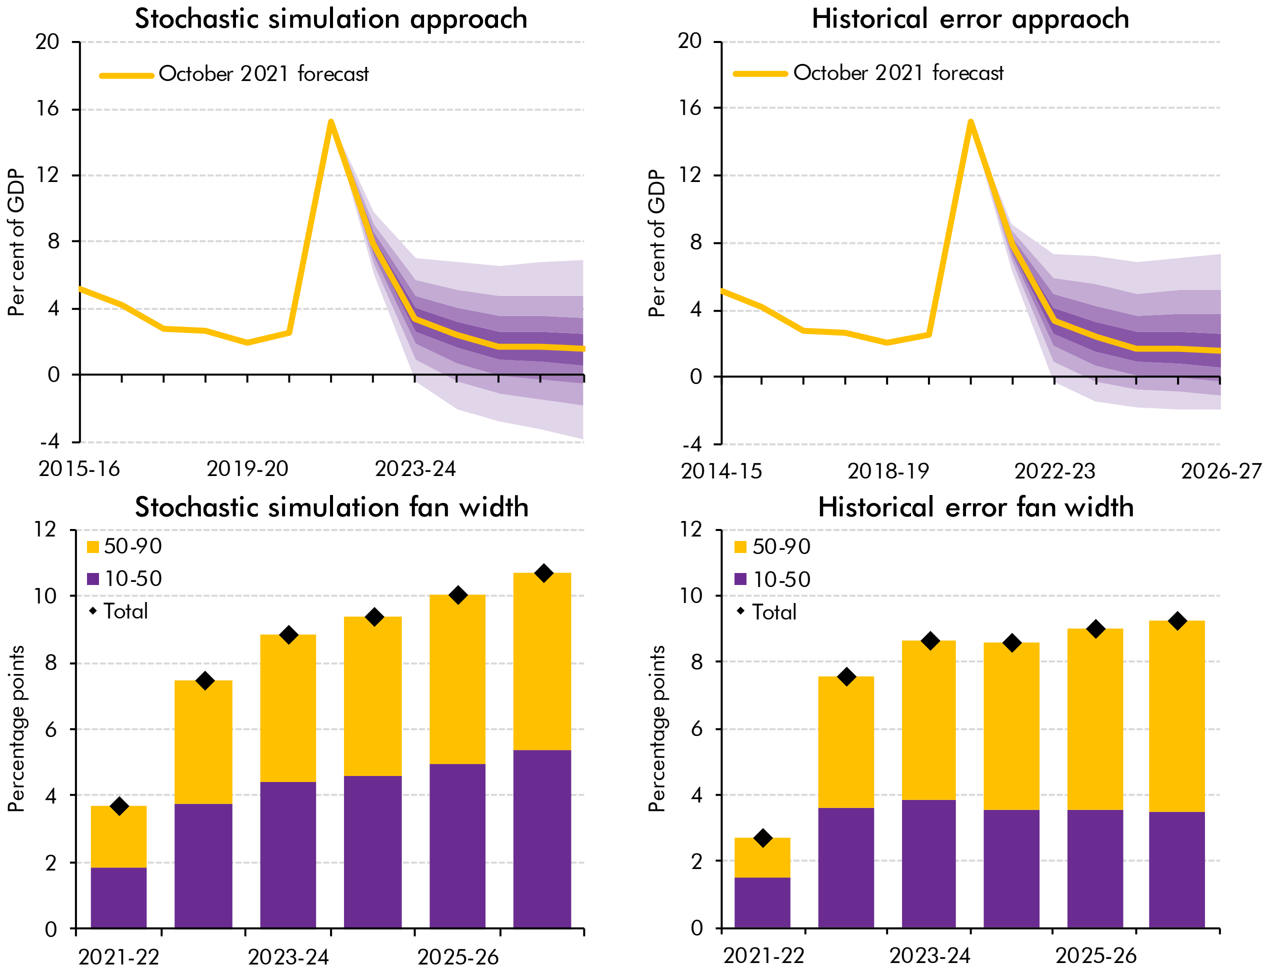

Source: ONS, OBR

4.7 Chart 4.5 shows a stochastic simulation fan chart for PSND, where the median again matches our October 2021 forecast (there is no historical error fan chart to compare it to, as noted earlier). The fan has an upward skew and by 2026-27 is remarkably wide, with the bottom of the fan nearly 40 percentage points below the median and the top nearly 60 percentage points above it. This is wider than many debt fans reported in the literature, probably reflecting the use of stabilising fiscal reaction functions that limit the movements in debt. It could also reflect the sample period used for drawing shocks, as most papers covered were written before the pandemic, so do not include the recent exceptionally large shocks (which would increase the upper tail of the fan), and most do not include a sample period back far enough to capture the high inflation of the 1970s (which could increase the lower tail of the fan).

4.8 The bottom panel of Chart 4.5 shows how the fan changes when some of the key judgements are varied. First, the sample period from which the shocks are drawn is very important. When the sample ends in the final quarter of 2019, so excluding the pandemic, the fan unsurprisingly narrows (by 23 percentage points in 2026-27). It also gives the fan a smaller upward skew. The effect of changing the sample period from which the shocks are drawn is explored further in paragraph 4.16. Second, another key judgement was the inclusion of SFAs. When these are excluded, the fan narrows (by 6 percentage points in 2026-27), and the upward skew is also moderated. Third, another key judgement was the inclusion of model uncertainty. When this is excluded, the fan narrows (by 29 percentage points in 2026-27) and the upward skew is, once again, moderated.

Chart 4.5: PSND fan chart using stochastic simulation approach

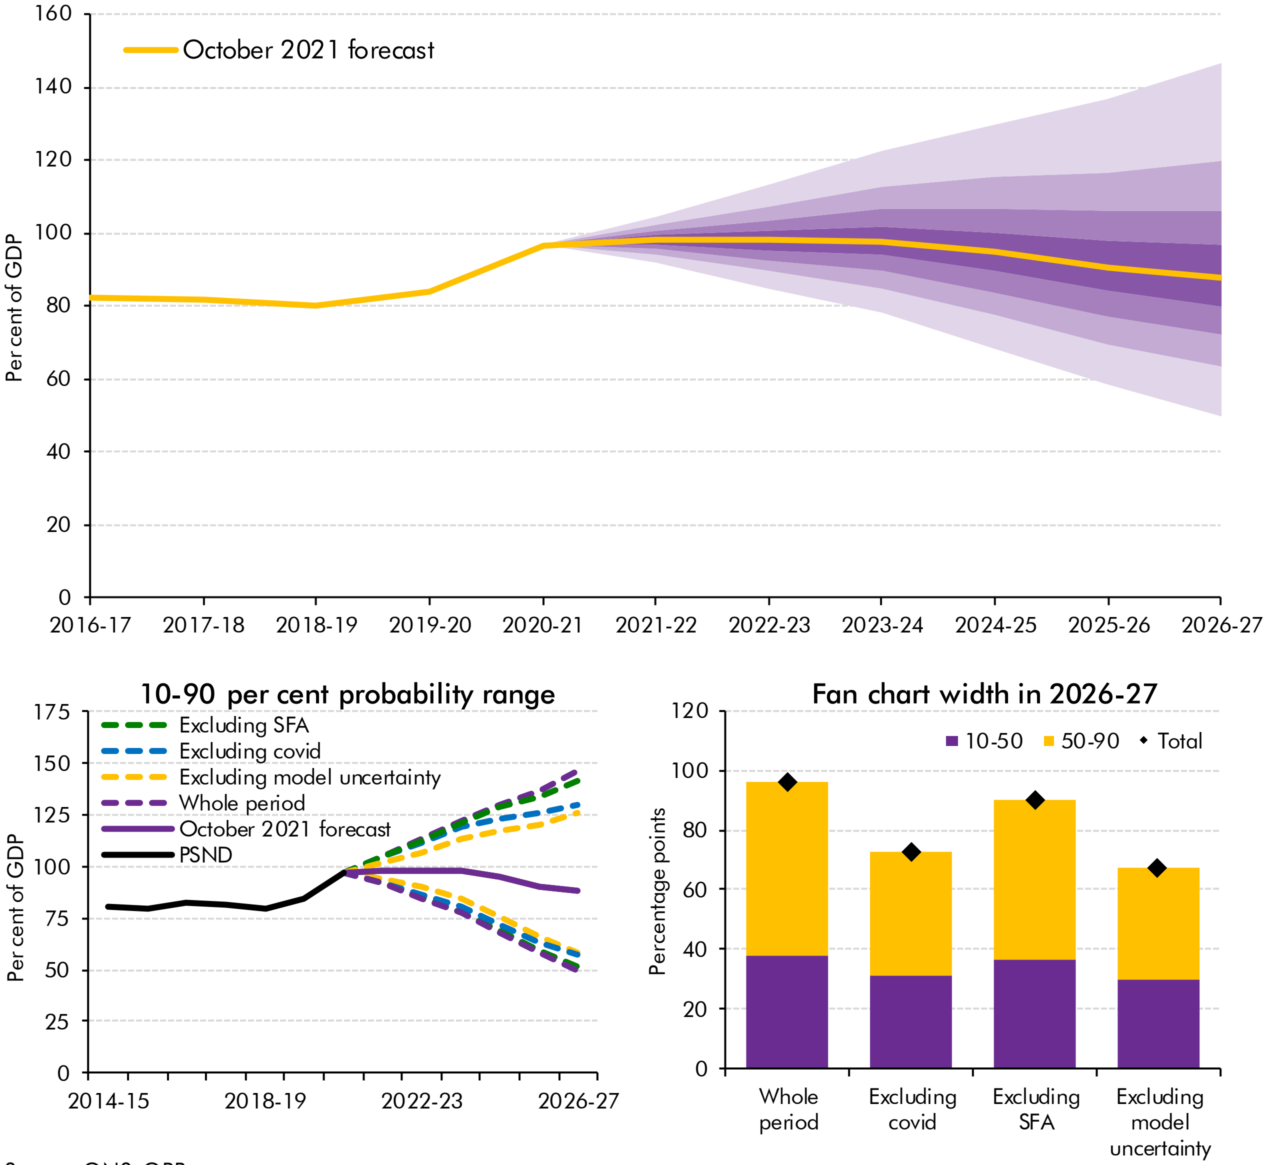

Source: ONS, OBR

Assessing fiscal rules

4.9 Although our model does not contain a structural fiscal reaction function, we can nevertheless use our approach to assess some potential fiscal criteria. For example, we can now estimate the probability of debt falling below a certain level in a given year, assuming that the fiscal reaction function embedded within the VAR remains unchanged. Chart 4.6 shows a fan chart for PSND excluding the effect of the Bank of England. This is the measure of debt that is used in the new fiscal mandate, which requires debt to be falling three years ahead (2024-25 in our October 2021 forecast). To calculate PSND excluding the Bank of England we assume that the size of the Bank's contribution to net debt is exogenous and therefore subtract our central forecast for it from each simulation of PSND. On this PSND measure, in 2024-25 (the current fiscal target year) there is a 20 per cent chance that debt will be between 80 and 90 per cent of GDP and an 80 per cent chance that it will be between 60 and almost 120 per cent of GDP. The width of the fan increases over time, with the 80 per cent interval widening from 26 per cent of GDP in 2022-23 to 92 per cent of GDP in 2026-27. The right-hand chart shows the probability that PSND as a share of GDP will be above a certain level (the probability that debt is below a certain level is therefore one hundred minus the probability shown in the chart). Each year the probability of attaining values significantly different to today's level increases.

Chart 4.6: PSND (excluding Bank of England)

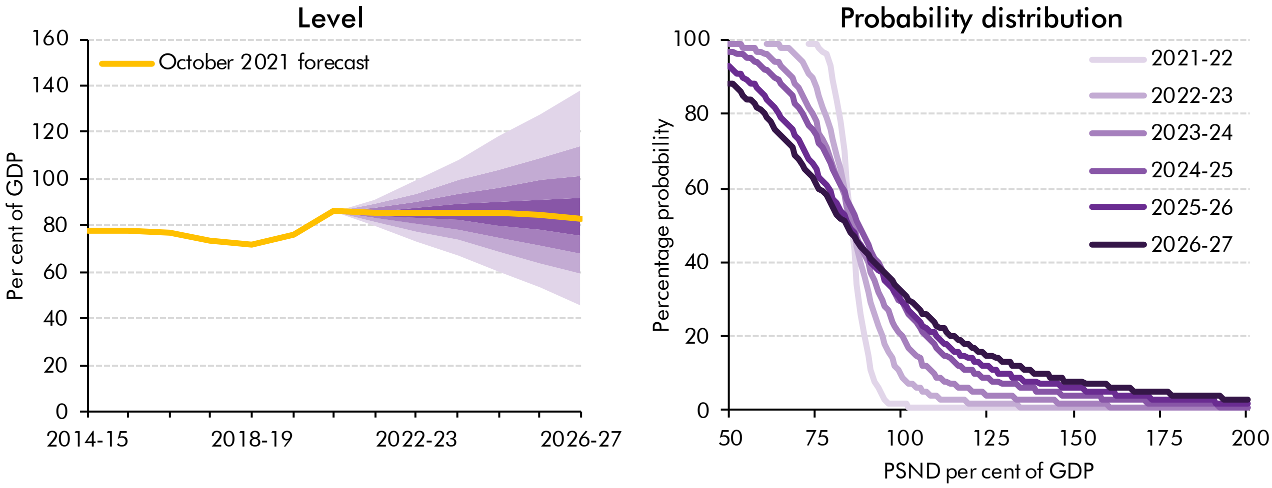

Source: OBR

Assessing fiscal rules individually

4.10 In our October 2021 *EFO* we looked at the probability of meeting the two new fiscal targets, which were announced alongside the Budget:

- to have public sector net debt excluding the Bank of England as a percentage of GDP falling by the third year of the rolling forecast period; and
- to balance the current budget by the third year of the rolling forecast period.

4.11 Both are rolling targets three years ahead, so each year the target moves forward one year. For our October 2021 forecast, the reference year was therefore 2024-25.

4.12 To assess the probability of debt falling in 2024-25, we can calculate the fraction of simulations in which debt falls in that year. This suggests a 54 per cent chance of debt falling and the fiscal mandate being met (Chart 4.7). This is the same probability suggested by the historical error approach.[24] The bands closest to the median of the fans are similar in both approaches so it is not surprising that they generate similar probabilities for meeting a target on which there is a modest headroom in our central forecast. The two fans differ more in the outer bands, where those of the stochastic simulation approach are somewhat wider.

Chart 4.7: Year-on-year change in PSND (excluding Bank of England) fan charts

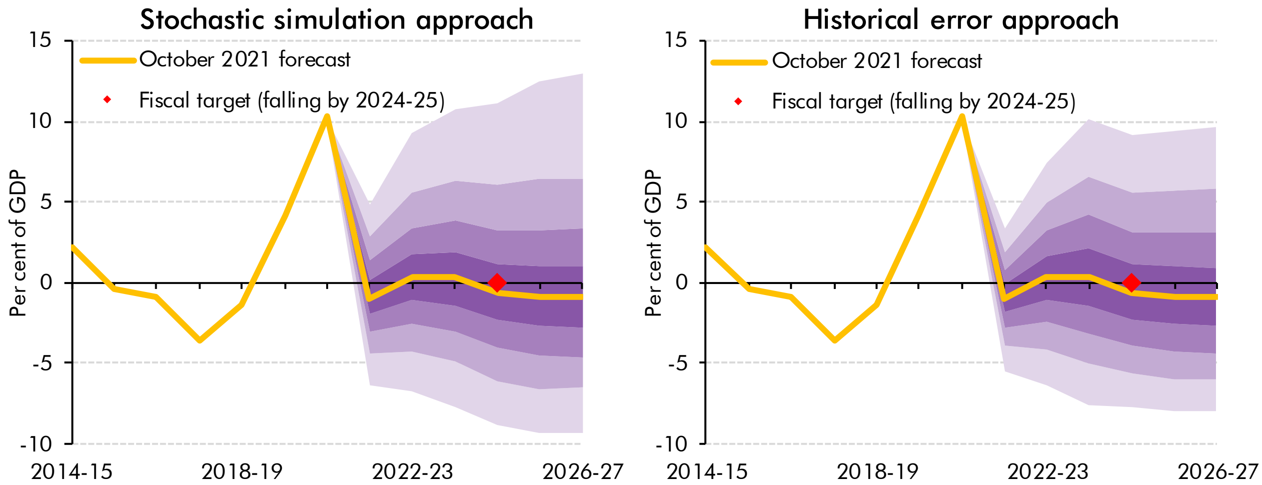

Source: ONS, OBR

4.13 Stochastic simulations can also be used to assess the probability of meeting the current budget balance target in 2024-25. To calculate the current balance, we assume that public sector net investment is exogenous and therefore subtract our central forecast for it from each simulation of PSNB. Chart 4.8 shows the probability fan around the current budget deficit, which gives a probability of this condition being met of 61 per cent. That is again the same probability as under the historical error approach (given the similarity of the central bands we would expect both approaches to give similar probabilities for a fiscal target met with a modest headroom in our central forecast). Again, the fans generated by the stochastic simulation approach are somewhat wider than those based on historical forecast errors, particularly in the near term. The current budget balance fan has less of a negative skew (or equivalently, there is less of a positive skew in the current deficit shown in Chart 4.8) in this approach. This reflects the inclusion of a different period of shocks, with the stochastic simulation approach including the very high levels of unanticipated inflation experienced in the 1970s. Drawing the shocks from 1998, as in the historical error fan chart, would produce a similar skew (though a wider fan).

[24] We can use the historical error approach to produce a fan chart for year-on-year changes in PSND by looking at past errors in individual years but as set out in paragraph 2.8 we cannot use it to produce a fan chart around the level of PSND across our five year forecast period.

Chart 4.8: Current budget deficit fan charts

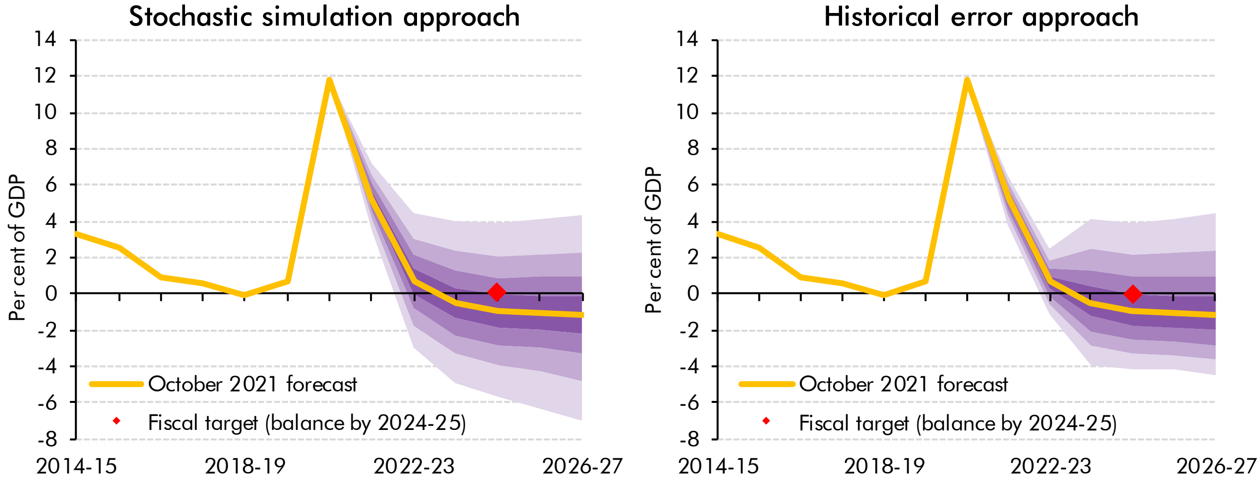

Source: ONS, OBR

Assessing fiscal rules jointly

4.14 We can also use the simulations to estimate the probability of meeting *both* these targets simultaneously in 2024-25 (the fraction of simulations where debt falls as a share of GDP *and* the current budget is in surplus). This is something we cannot do with the historical errors approach as we do not have a joint distribution of errors. Chart 4.9 shows a scatter plot of the different simulation values for the current budget and change in debt in 2024-25, where the darkness of the area signifies a concentration of outcomes. The bottom two quadrants show instances where PSND is falling (in 54 per cent of simulations), the two quadrants on the right show instances where the current budget in surplus (61 per cent of simulations), while the bottom right quadrant shows instances where both targets are met at the same time (just 40 per cent of simulations). This shows that while both targets are met in our central October 2021 forecast (the red diamond), on average across all the simulated futures there are more occasions where one or both targets are missed than there are where both are met together.

Chart 4.9: 10,000 simulations of current balance and change in PSND in 2024-25

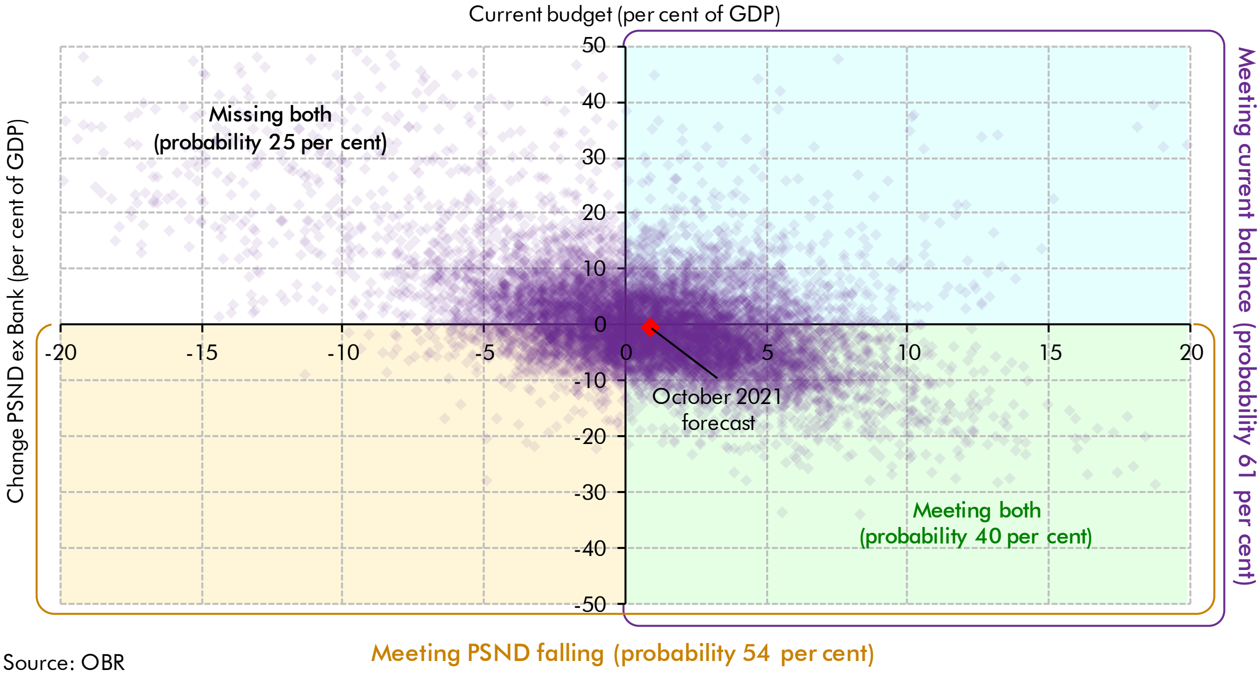

Source: OBR

4.15 Chart 4.9 also shows that while most simulation values in the target year are relatively close together there are some large outliers – many in the top left quadrant. As Table 4.1 shows, 4.1 per cent of the 10,000 simulations have PSND rising by over 20 per cent, while 2.4 per cent have a current deficit of greater than 10 per cent, with 1.6 per cent of the simulations meeting both of these conditions (the large outliers in the top left quadrant of Chart 4.9). That is similar to the proportion of pandemic-affected years in our sample, i.e. 1 in 65 or 1.5 per cent. There are fewer outliers in the other quadrants of the chart.

Table 4.1: Proportion of outlier simulation values in 2024-25 (per cent)

	PSND change	Current budget	Joint
	>20 per cent	<-10 per cent	Top left quadrant
Proportion	4.1	2.4	1.6
	>20 per cent	>10 per cent	Top right quadrant
Proportion	4.1	1.7	0.2
	<-20 per cent	<-10 per cent	Bottom left quadrant
Proportion	0.8	2.4	0.0
	<-20 per cent	>10 per cent	Bottom right quadrant
Proportion	0.8	1.7	0.3

Sensitivity of fiscal rules assessment to the period shocks are drawn over

4.16 The width and skew of the fans are affected by the sample period from which the shocks are drawn, which will also affect the probability of meeting fiscal rules. But as we align the median of the fan chart with our central forecast, the period shocks are sampled from will not affect whether the rule is met or missed (which is determined in our central forecast assuming the rule is set to target a 50 per cent probability of being met). Chart 4.10 shows the width of the fans for different shock sample periods. When the pandemic period is excluded (i.e. drawing from 1956 to the final quarter of 2019), the width of the fan is 21 percentage points narrower in 2026-27 than when the shocks are drawn up to the second quarter of 2021. There is also less of an upward skew. The width of the fan chart narrows further when the period is set to the 'Great Moderation', the particularly benign period running from 1993 to 2006. Drawing only over the Great Moderation, the fan chart is 29 percentage points narrower by 2026-27 than drawing over the whole period. It is also the least upwardly skewed of the fan charts. Finally, drawing over the period of 1998 to the second quarter of 2021 (the period used for the historical error fan charts in the section above) produces a noticeably wider fan than using the whole period. This is because it includes both the financial crisis and the pandemic, but the sample period is much shorter, so those two very large shocks get drawn more frequently. The fan it produces also has the largest upward skew of the sample periods shown. This illustrates the importance of assuming an appropriate comparator period for the calibration of the fan chart.

Chart 4.10: PSND fan charts for different shock sample periods

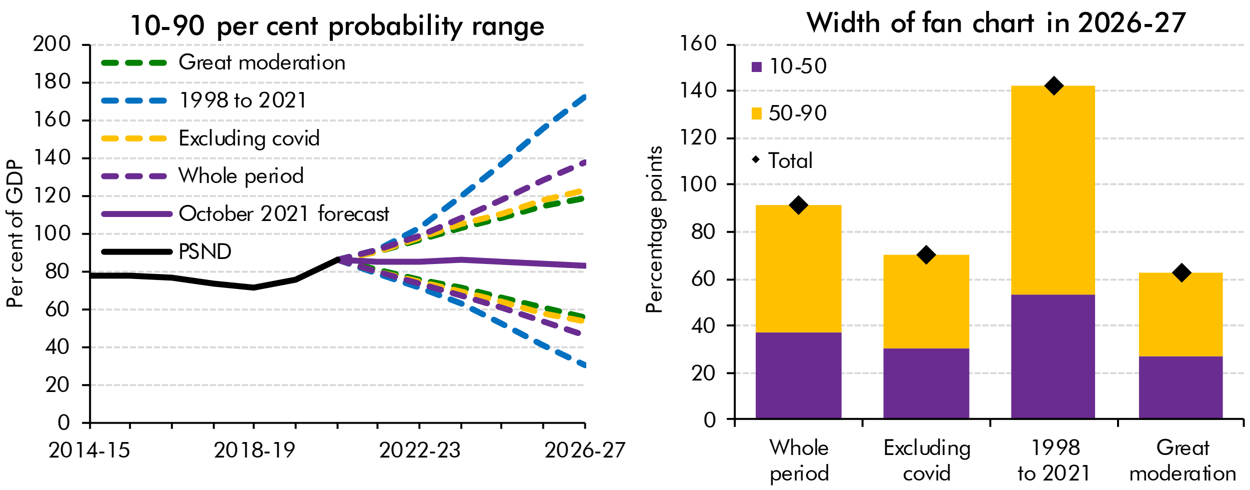

Source: ONS, OBR

4.17 The sample period over which shocks are drawn has the smallest effects on the bands closest to the median. Consequently, the probability of meeting a fiscal rule with a modest headroom in our central forecast is not that sensitive to the sample period chosen for the shocks. Table 4.2 shows the probability of meeting fiscal targets for the October 2021 forecast for different shock sample periods, with the probability of meeting the debt-falling target ranging only very modestly from 54 to 56 per cent, and the current budget balance target ranging from 60 to 63 per cent.

Table 4.2: Probability of meeting fiscal targets using different shock sample periods

	Debt falling	Current balance	Joint
Whole period	54	61	40
Excluding pandemic	54	62	39
1998 to 2021	54	60	42
Great Moderation	56	63	41

Possible future improvements

4.18 We believe our stochastic simulation approach is now sufficiently developed for us to start using it at our next *EFO* in the Spring as our primary method for calibrating fan charts and assessing fiscal rules. But we expect to continue to refine our approach in the future. Ahead of the next forecast we will re-estimate the model equations for the latest ONS Blue Book data, and when necessary update them every year for the autumn Blue Book release. There are several further, more complicated, methodological improvements that we may investigate over a longer timescale, such as:

- Extending the VAR, for instance by including lags of debt-to-GDP as exogenous variables to capture effects on interest rates and the primary deficit. Replacing the VAR equation for the primary deficit with a structural fiscal reaction function would also allow us to assess the ramifications of different fiscal rules if required by government;

- Improving the modelling of SFAs;

- Modelling the elements of the simulations that we have assumed to be exogenous and therefore used our central forecast values for, such as the size of the APF.

A Data sources and diagnostics

Variable descriptions and sources

Table A.1: Data and sources

Series	Variable symbol	ONS/Bank code	Frequency used	Note
Real GDP	g (growth)	ABMI	Quarterly	
GDP deflator	π (inflation)	(YBHA/ABMI)	Quarterly	
Nominal GDP		YBHA (SA) KLS2 (NSA)	Quarterly	
Primary deficit	PD	-(J5II-JW2L+JW2M+JW2P)	Quarterly	Series seasonally adjusted
Bank Rate	BR	IUQABEDR	Quarterly	
Effective interest rate (ex APF)	EIRx	(JW2P-JW2L+JW2M-MDD7)/(HF6W-MEX2)	Quarterly	Series seasonally adjusted
Effective interest rate	EIR	(JW2P-JW2L+JW2M)/HF6W	Quarterly	Series seasonally adjusted
Public sector net debt	PSND	HF6W	Annual/ Quarterly	Pre-1974-75 annual data from Bank's "A millennium of macroeconomic data". Not available quarterly pre-1993. NCR used to create quarterly back series.
Public sector net debt (ex APF)	PSNDx	HF6W-MEX2	Annual/ Quarterly	Series seasonally adjusted
Public sector net debt (ex Bank of England)	PSNDxB	CPPH	Annual	
Public sector net borrowing	PSNB	-J5II	Annual	
Public sector net investment	PSNI	-JW2Z	Annual	
Current balance	CB	JW2T	Annual	
Asset purchase facility	APF	MEX2	Annual/ Quarterly	
Stock-flow adjustment	SFA		Annual	Difference in calculated change in PSND and actual change.

Diagnostics

Table A.2: Unit root tests

Variable	Phillips-Perron test statistic	P-values
Real GDP growth	-15.46	0.000**
Inflation	-10.44	0.000**
Bank Rate	-1.69	0.434
Effective interest rate (ex APF)	-1.46	0.551
Primary balance (per cent GDP)	-2.88	0.049*

Memo: Null hypothesis is variable has a unit root.
*Reject at 5 per cent significance level.
** Reject at 1 per cent significance level.

Table A.3: VAR serial correlation test

Lag	VAR 4-lags		VAR 5-lags (specification used)	
	Rao F-stat	P-value	Rao F-stat	P-value
1	2.04	0.002**	0.94	0.553
2	1.81	0.001**	0.97	0.534
3	1.63	0.001**	1.11	0.249
4	1.47	0.003**	1.15	0.152
5	1.36	0.008**	1.20	0.073
6			1.19	0.072

Memo: Null hypothesis is no serial correlation at lags 1 to X.
* Reject at 5 per cent significance level.
**Reject at 1 per cent significance level.

Chart A.1: VAR stationarity test (inverse root circle)

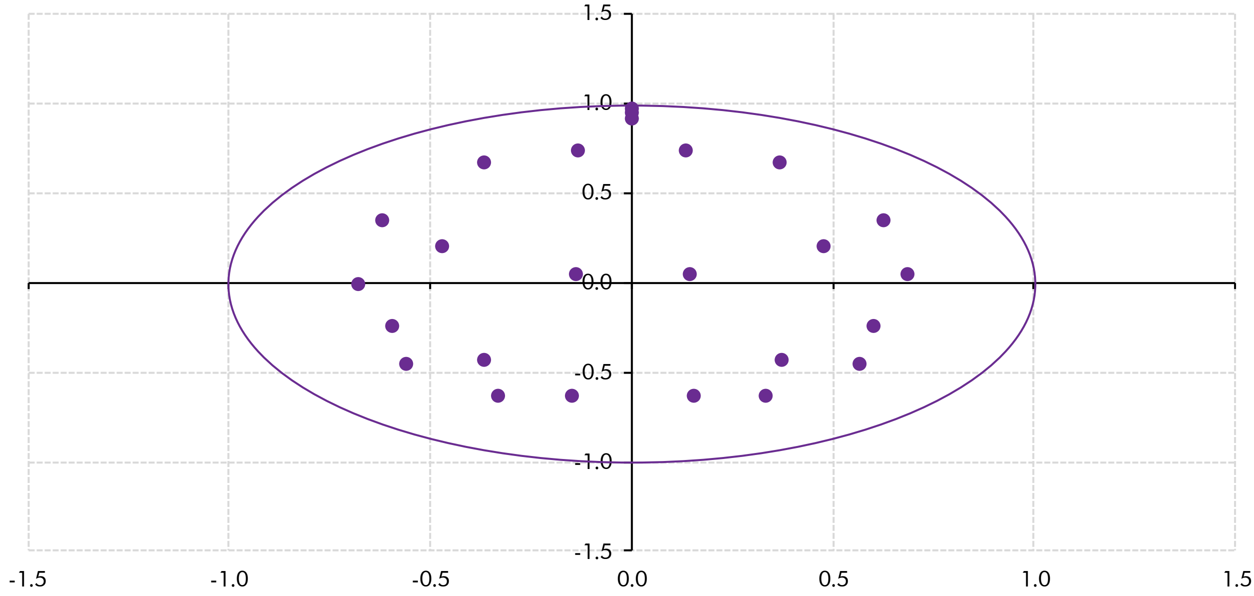

Note: If no roots are outside the unit circle then the VAR satisfies the stability condition.
Source: OBR

Chart A.2: VAR residuals (third quarter of 1956 to fourth quarter of 2019)

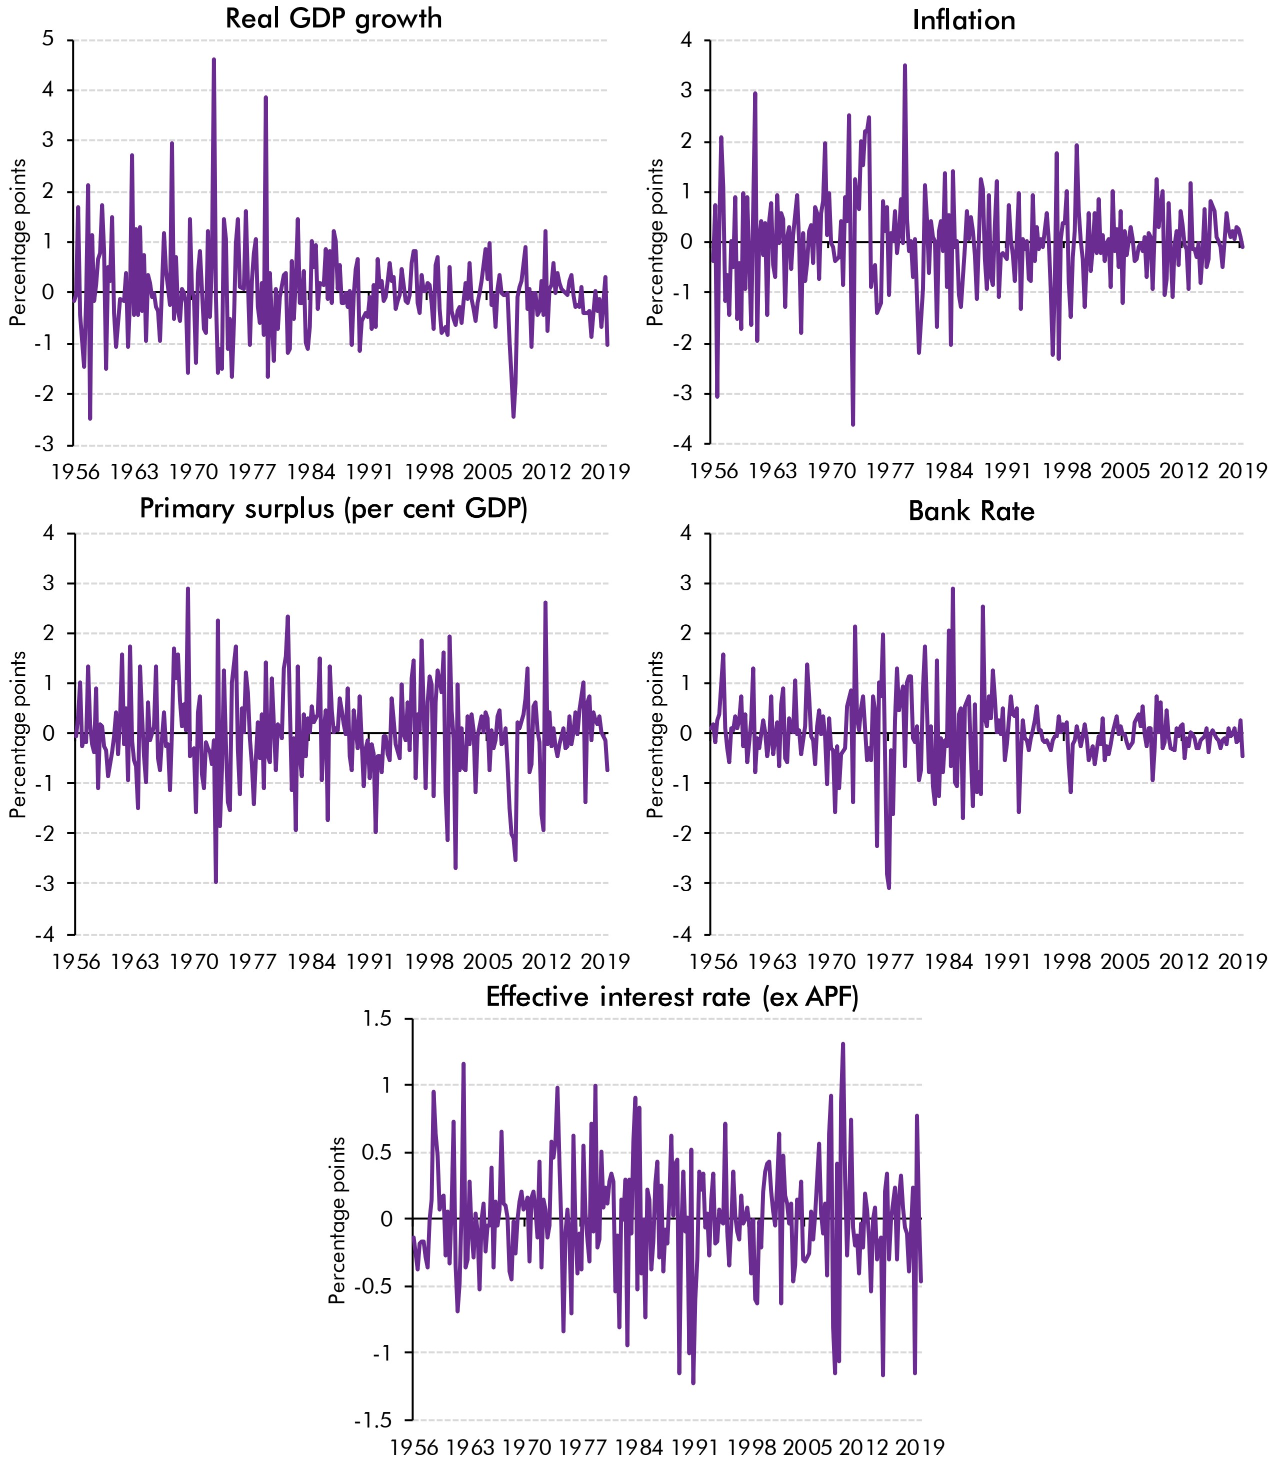

Source: OBR

Table A.4: VAR normality of residuals test

	Chi-squared statistic			P-value		
Component	Skewness	Kurtosis	Jarque-Bera	Skewness	Kurtosis	Jarque-Bera
1	48.94	260.23	309.16	0.000**	0.000**	0.000**
2	0.05	25.25	25.30	0.825	0.000**	0.000**
3	0.67	7.17	7.84	0.413	0.007**	0.020*
4	0.84	117.97	118.80	0.361	0.000**	0.000**
5	1.10	7.81	8.92	0.293	0.005**	0.012**
Joint	51.59	418.43	470.02	0.000**	0.000**	0.000**

Memo: Null hypothesis is residuals follow a multivariate normal distribution.

* Reject at 5 per cent significance level.

** Reject at 1 per cent significance level.

Table A.5: Stock-flow adjustment equation results

Variable	Coefficient	Standard error	t-statistic	P-value
Constant	-0.19	0.164	-1.17	0.25
GDP dummy	3.72	0.536	6.95	0.00**

Memo: Adjusted R^2 *0.43. GDP dummy takes value 1 if real GDP falls, 0 otherwise.*

**Significant at 1 per cent level.

References

Banerjee, N., and Abhiman, D., (2011), *"Fan chart: methodology and its application to inflation forecasting in India."* RBI working paper series.

Berti, K. (2013), "*Stochastic public debt projections using the historical variance-covariance matrix approach for EU countries*" Directorate General Economic and Financial Affairs (DG ECFIN), European Commission.

Bouabdallah, O., et al. (2017), *"Debt sustainability analysis for euro area sovereigns: a methodological framework."* ECB Occasional Paper, (185).

Britton, E., Fisher, P., & Whitley, J. (1998), *"The Inflation Report projections: understanding the fan chart." The Bank of England*

Carton, B., & Fouejieu, A. (2020), "*Assessing Dutch Fiscal and Debt Sustainability*." IMF working Paper No. 2020/269

Casey, E., & Purdue, D. (2021), "*Maq: A Macro-Fiscal Model for Ireland*." Irish Fiscal Advisory Council Working Paper Series No. 13.

Cherif, R., & Hasanov, F. (2012), "*Public debt dynamics: The effects of austerity, inflation, and growth shocks*" IMF Working Paper No. 12/230.

Cuerpo, C., & Ramos, J. M. (2015), *"Spanish Public Debt Sustainability Analysis" Hacienda Publica Espanola*, (215), 95.

European Commission (2021), "*Debt sustainability monitor 2020*" European Commission Institutional Paper 143.

IMF (2011), *"Modernizing the framework for fiscal policy and public debt sustainability analysis"* IMF *Policy Papers*, *2011*(034)

IMF (2021), *"Review of the Debt Sustainability Framework for Market Access Countries"*, IMF Front Matter, *Policy Papers*, *2021*(003)

Irish Fiscal Council (2021), *"Fiscal Assessment Report"*

Jaramillo, L., Mulas-Granados, C., and Kimani, E. (2016), "*The blind side of public debt spikes*", IMF working paper number 16/202.

Johnson, N. L., Kotz, S., & Balakrishnan, N. (1994), "*Lognormal distributions. Continuous univariate distributions*", *1*, 207-27.

Lukkezen, J., & Rojas-Romagosa, H. (2012), "*When is debt sustainable?*." CPB Netherlands Bureau for Economic Policy Analysis.

Medeiros, J. (2012), "*Stochastic debt simulation using VAR models and a panel fiscal reaction function–results for a selected number of countries*" Directorate General Economic and Financial Affairs (DG ECFIN), European Commission No. 459.

OECD (2020), "*OECD independent fiscal institutions review: Office for Budget Responsibility of the United Kingdom*"

Office for Budget Responsibility (2012), *"How we present uncertainty",* OBR Briefing Paper No. 4

The Network of EU Independent Fiscal Institutions (2021), "*The role of the Independent Fiscal Institutions in assessing the sustainability of high public debt in the post-Covid era*"

Ufficio Parlamentare di Bilancio (2019), "*2020 Budgetary Policy Report*"

Výškrabka, M. (2016), "*Stochastic forecast of the Slovak public debt*", MFSR.