

UK Budget Preview 1: What are fiscal rules, why do we have them, and could they be made better?

With the UK Budget approaching, speculation is mounting that Chancellor Rachel Reeves is considering changing the fiscal rules that the Office for Budget Responsibility (OBR) uses to assess whether the Government is meeting its budgetary targets. So it's worth thinking about how they have come about, and to what extent the rumoured changes are (i) sensible and (ii) likely to make a difference.

Why do we have fiscal rules?

Macroeconomists worry that fiscal policy suffers from what they call a 'time inconsistency' problem. This means that what governments might see as their better option today might not be optimal in the long run. For example, a government close to an election might cut taxes or increase public spending in order to be re-elected, even if that causes public debt to increase faster or if it contributes to higher inflation by increasing aggregate demand at a time when such a stimulus isn't warranted.

To some extent, governments also suffer from the very human 'time-preference' problem: we'd rather postpone painful decisions if we can, but if we do so indefinitely, it will eventually catch up with us. Why take unpopular decisions that make you less likely to hold on to power, only for the next government to benefit from them, as the Conservative Party found under John Major?¹

Even if a looser fiscal policy might feel the path of least resistance at any point in time, persistent macroeconomic imbalances or unexpected increases in government borrowing with no plan for re-balancing the budget often do have very large and expensive consequences for the economy as a whole. There is extensive academic literature on the 'deficit bias', and how it can lead to poor economic outcomes, such as debt crises or slowdowns in growth.

But there are also solutions, which help navigate these issues to some extent.² As with all mechanisms designed to solve time inconsistency problems, they must tie the decision-makers' hands to some extent by removing discretion – otherwise, they wouldn't be an effective solution. Options include specific rules about how large the deficit or the level of public debt can be or having some sort of external assessment of whether these rules are being met. The UK has both, though they were introduced at different times.

¹ Sanders, David. 1999. "Conservative incompetence, Labour responsibility and the feelgood factor: why the economy failed to save the Conservatives in 1997." *Electoral Studies*, Volume 18, Issue 2, June 1999, Pages 251-270.

² Krogstrup, S., Wyplosz, C. (2009). "Dealing with the Deficit Bias: Principles and Policies." In: Ayuso-i-Casals, J., Deroose, S., Flores, E., Moulin, L. (eds) *Policy Instruments for Sound Fiscal Policies*. Finance and Capital Markets Series. Palgrave Macmillan, London, Pages 23-50.

A history of fiscal rules in the UK

Formal fiscal rules are a relatively recent innovation in the grand scheme of UK fiscal history, and date back to Gordon Brown's appointment as Chancellor of the Exchequer in 1997.³ Nevertheless, implicit or informal rules have guided policy for a long time. For example, prior to the Second World War, and with the exception of periods of international conflict, the government tried to achieve a balanced budget, and would adjust tax rates every year to achieve that. During the 19th century, it also attempted to reduce debt by a specified amount by having enough funds to repay it in net terms – something that we would today define as a primary surplus rule.

When Gordon Brown took office, two rules were introduced:

- The 'golden rule', which stated that the current budget should be in balance over the economic cycle; and
- The 'sustainable investment rule', which stated that the debt-to-GDP ratio should not average more than 40% over the economic cycle.

Apart from being misnomers – the golden rule in economics has to do with optimal economy-wide saving rather than the public finances, and the 'sustainable investment rule' had nothing directly to do with investment – these were also fairly murky rules. The economic cycle isn't easy to date, with vintages of data and revisions changing our understanding of when the high and low points of GDP over a cycle took place. It also didn't help that the UK doesn't have a business cycle dating committee like the one in place in the US through the National Bureau of Economic Research. Looking back through the 1997 to 2009 budgets, one can find an ever-changing landscape of economic cycles, which the Treasury got to define to its own benefit.

For all their problems, the stability of these rules is quite remarkable relative to what came next. Some of this was due to the Great Financial Crisis – faced with such a large shock, the level at which rules were set looked impossible to achieve for a long time and attempting to do so would be counterproductive.

But there is a feeling that the first formal change was like letting the genie out of the bottle, and subsequent Chancellors have felt empowered to change them much more often. Contrast this with the 2% inflation rule that the Bank of England has had for the whole of that same period: even when it missed it by a large amount, that has been maintained and the Bank has instead signalled what it was doing to bring it back into balance. But perhaps that is the difference between a political position like that of Chancellor and the operationally independent Governor of the Bank.

Since 2009, we are on our eighth set of fiscal rules, and Rachel Reeves has promised yet another one. All of these have generally measured sensible things⁴ – the current budget deficit, the overall deficit, public sector net debt (including or excluding the Bank of England), the cyclically adjusted deficit –

³ <https://www.instituteforgovernment.org.uk/explainer/fiscal-rules-history> (accessed 10 October 2024).

⁴ There is much less of a case for positive assessments to be made about the public sector net investment rule – which was simply a bizarre way of combining with the current budget rule to create an overall borrowing rule. The welfare cap was also a fiscal rule for a while, but it was essentially a political gimmick that did not cap welfare spending as a whole in any meaningful way (it excluded the state pension, the single largest benefit paid by the UK Government) and only served to bear down on in-work benefits.

but the inability to stick to a set of rules has undoubtedly led to a loss of some of the confidence in their worth. After all, why should the next set be the right one if it changes every couple of years?

With the OBR into being in 2010, responsibility for assessing compliance with the fiscal rules was transferred to the nascent institution. This meant an improvement in transparency and in stringency of the rules, as no longer would some of the judgements made internally in the Treasury about the economic cycle or trend growth be allowed to sway the likelihood of the rules being met.

Before coming into Number 11, Rachel Reeves set out what she intended to install as her fiscal rules. In her Mais Lecture, delivered in March 2024, the then-Shadow Chancellor said that she would return to having the current budget in balance, and that debt would be falling in the final year of the forecast.

[As we said before the election](#), the difference this will make in practice is immaterial – the debt rule is the binding constraint, and therefore if it remains in place (even if it's the supplementary target), it will be what determines how much the Government can spend without the OBR deeming it to be failing to meet the fiscal rules. Removing the Bank of England/Central Government split might give the Chancellor some room for manoeuvre, but it would not be a massive game-changer.

'Fiscal rule nutters' and Goodhart's law

When inflation targeting was becoming the norm, Mervyn King – later to become Governor of the Bank of England – worried about the possibility for an 'inflation nutter' – someone who only cared about bringing inflation to target without considering employment or GDP – being in charge of the central bank and the damage that could do to the economy.⁵

That has not come to pass, and central bankers have been much more flexible in their approach, being allowed the time and the discretion to change the speed at which inflation is adjusted back to target. This means that the inflation target has on many occasions been missed, but it still proves a useful anchor for where the Bank of England is headed. Crucially, the Bank has no qualms about forecasting that the target will not be hit in the forecast horizon, and does not feel compelled to set out an overly aggressive course of action just for the sake of hitting it.⁶


Contrast this with the fiscal rules. Since the OBR has been introduced and rigour enhanced in the scrutiny of the rules in 2010, they have never been forecast not to be met at a fiscal event⁷ (see chart 1).

⁵ King, M. 1997. "Changes in UK monetary policy: Rules and discretion in practice." *Journal of Monetary Economics*, Volume 39, Issue 1, June 1997, Pages 81-97.

⁶ See for example the [May 2023](#) and [February 2024](#) Monetary Policy Reports.

⁷ This excludes the welfare cap, which was never a real constraint on overall fiscal policy.

Chart 1: Headroom against meeting the main fiscal rule since 2010



Source: OBR. Values are calculated in terms of 2028-29 GDP forecasts for comparability.

On its face, this seems like a great success. And the 2022 mini-budget – not an official OBR forecast, but a set of policy announcements which clearly implied the fiscal rules would not be met – would seem to vindicate the success of the fiscal rules. But leaving aside the mini-budget violated all assumptions of what sustainable fiscal policy and market chaos ensued, the apparent success in continually meeting every set of fiscal rules is mitigated by three factors:

1. A number of the rules targeted particular dates at which the deficit was meant to be of a certain magnitude. These rules are essentially hostages to fortune – if there is an economic shock, they are unlikely to be met. In all these cases, by the time the target date came around, they have been ditched as rules – not unrelatedly because they were unlikely to be met in practice.
2. That chopping and changing of rules coincided with large changes in economic conditions, which themselves largely caused the rules in place not to be achievable. This means that instead of confronting the fact that the rule was not likely to be achieved, the Government decided to abandon it and replace it with a more achievable rule. This can make sense and make rules more meaningful, and can also prevent governments from being hamstrung by a bad rule – for example, the so-called ‘debt brake’ in Germany, which has been extremely damaging to growth but is politically untouchable. But the fact that these changes are instigated by the Government itself rather than an independent authority creates an

appearance of moving the goalposts to suit the results rather than complying with spirit of the rules.

3. The Treasury's fine-tuning of fiscal policy to just about comply with the letter of the rules at all times has created an institutional culture in which being assessed as not complying with the rules in place is not an option.

This fine-tuning has given rise to the Treasury behaving as 'fiscal rule nutters', to borrow Mervyn King's turn of phrase. Whatever the starting point for economic conditions and whatever policy is in the short-term, the Treasury's ultimate aim when setting medium-term policy is to meet the fiscal rules – in its current form, debt excluding the Bank of England to fall from year 4 to year 5.

The incentives this creates are extremely poor and damaging. For example, in the 2022 Autumn Statement, the fiscal rules were met on the basis of 0.7% real-terms cuts to unprotected government departments in the period after the Spending Review, as well as holding capital budgets flat in cash terms – i.e. cutting them in real terms and as a share of GDP.

Why was that chosen as the path? Clearly there are political options at play, but there are very understandable if ultimately damaging practical reasons at play. Raising taxes creates real losers, and people generally tend to pay the same taxes year-on-year – and therefore they know immediately that they will have to pay more in future. Departmental budgets, however – especially when spending has yet to be disaggregated across areas – are much less concrete, and therefore the path of least resistance. This is a well-trodden path, and one which has meant that policy upon policy decision over the last couple of years has been predicated on these quite draconian departmental spending plans.

In a very real sense, the 'fiscal rule nutter' tendency exhibited by the Treasury is an example of Goodhart's law – that when a measure becomes a target, it ceases to be a good measure. The obsession with a particular fiscal rule being met means that all matter of damaging but less immediately unpalatable or visible options come into play in order to achieve the desired metric.

So are fiscal rules actually needed, and what would be a better way forward?

The obsession with meeting fiscal rules is to some extent the inevitable corollary of the spelling out and condensing of fiscal rules into one or two metrics, as well as the regular and required assessment of them. Given official primacy, these rules become the focus of outside observers in bringing the Government to account, and mean that it is politically very difficult to break with the single metric framework.

This criticism of 'fiscal rule nutterism' is therefore not directly aimed at persons or institutions, but rather at the system of incentives in which they operate. They are hamstrung by it, unable to change it meaningfully and yet – in many cases – fully aware of the damaging consequences it has. And once fiscal rules have been introduced, the signal that getting rid of them altogether would send would be very costly, and potentially lead to market disruption – again, the 2022 mini-budget is a cautionary tale.

Not to mention that fiscal sustainability is important in and of itself. A credible path for tax and spending is crucial for maintaining market confidence; on the other hand, a clearly unsustainable path is likely to cause the holders of government debt to worry about the certainty of repayment and the extent to which they will be protected from consequences like higher inflation.

It's important to go back to basics and consider what the mechanism for a crisis of confidence manifesting itself. Ultimately, the biggest constraint a government faces when trying to finance itself is whether it can find enough people to buy its issuance of debt and at what cost. This is why so many fiscal rules concern themselves with debt and interest costs, and with good reason.

This is not dissimilar from the old business adage that 'cash is king': it's all well and good to have accounting profits, but if they're not generating cashflow, a business might still go under. Likewise, if a government needs to spend actual money, the (accrued) deficit position can be pretty misleading – it's what the Treasury calls the net cash requirement that ultimately matters.

This is another reason why some of the clamour for broader rules such as targets for public sector net financial liabilities and public sector net worth should give us pause. This doesn't mean that they don't measure useful or meaningful concepts. But their relationship with financial market metrics is much looser than that of net debt, and it's worth pondering whether targeting them with no reference to debt is the ideal way of bounding fiscal decisions.

In fact, the main takeaway should be that a single metric is not the best way to ensure both fiscal sustainability and good decision-making. It becomes too politically charged and governments feel unable to bear the political cost of not meeting a single target that is meant to encapsulate their fiscal responsibility or lack thereof.

Instead, a more grown-up way forward would be to instruct the OBR to assess a broader set of metrics. It already does so for most of them, compiling them in a dashboard – [table 5.3](#) of the [most recent Economic and Fiscal Outlook](#). These include public sector net debt (including and excluding the Bank of England), public sector net financial liabilities, public sector net worth and debt interest costs – both as a share of revenues and as a share of GDP. Combining this with the cyclically adjusted budget deficit, overall deficit, primary deficit and monitoring them relative to the beginning of the forecast rather than from year 4 to year 5 would provide a broad set of metrics that could then be interpreted as a whole.

Of course, a set of indicators not as simple to interpret as a single number, and it may feel like it hinders public accountability. But the simplicity is deceptive in the single metric case. It makes it look like the metric embodies a complex set of judgements. But metrics are fallible and have limitations, and relying on a single one is being hostage to fortune.

[As Einstein might or might not have said](#) in quite these words, 'everything should be made as simple as possible but no simpler.' Fiscal sustainability is a complex matter. It's one where we can only really know after the fact whether it was fulfilled or not – and even then, there can be a lot of debate. Ultimately what the Government should be trying to achieve is to make decisions on the basis of sound macroeconomic management, as well as sending a signal to financial markets that its decisions are considered and compatible with fiscal sustainability – rather than making fanciful assumptions about raising fuel duty for the first time in a decade and a half or about bearing down on departmental spending without having to consider what delivering such a change would mean for public services.

The UK's budget process also encourages tinkering and fine-tuning – and that needs to change urgently

Prior to the OBR's introduction, the Treasury's process of putting the Budget together involved endless tinkering and fine-tuning – officials still today tell tales of forecast rounds being added after the process was meant to have been closed, re-jigging the fiscal package to ensure the right number was reached. There is no suggestion of impropriety, merely that it caused the budget to be re-written at the very last minute, possibly after more back-and-forth than necessary and desirable.

The OBR brought some much-needed stability and structure to the process. The OBR occupies an uncommon place in budgeting, especially when compared with other countries' similar bodies. It is fully embedded into the forecast process: the Treasury must use its forecasts when presenting a budget, and that means a much more frequent flow of information between the OBR and Number 11.

There are some significant advantages of this setup. The OBR is able to significantly challenge the UK Government's narrative and hold it to account on the day the budget is published. In other jurisdictions, fiscal watchdog assessments tend to follow the presentation of the budget – by the time they report back, the Government has had the opportunity to build the narrative of the fiscal event. So the OBR's position in the centre of the process really does increase accountability, and it's a very valuable advantage over alternative institutional designs.

But the Treasury still yields significant influence over the timetable – more so than would be desirable, and that control has increased over time. The typical OBR forecast used to look like as follows:

- Three rounds of iterative 'pre-measures' forecasts, in which modelling changes and new data are incorporated;
- Notification of 'major measures' from the Treasury – essentially setting out the size of the package, even if details can then be ironed out;
- A fourth round, including the effects of those major measures on macroeconomic variables;
- A final '[scorecard](#)', including all the measures that the Treasury wanted to include in the forecast; and
- A final forecast, essentially Round 4 plus final scorecard, returned to the Treasury the Friday before the Budget.

This already allowed the Treasury plenty of opportunity for fine-tuning, but crucially it restricted its ability to attempt to achieve overly precise headroom numbers by not iterating them. Since then, the Treasury has succeeded in adding an iteration of the 'major measures' notification – essentially giving it an extra go at achieving the 'desired' amount of headroom. For this budget, that even came at the [cost of the third 'pre-measures' round](#), suggesting an institution even less shy of using its power.

This might not seem like a crucial change, but it does indicate a disappointing direction of travel. It feeds the headroom obsession that appears to have grown ever more prominent in recent years. When looking five years down the line, many of these small differences are meaningless – we can easily differentiate between the 0 and 1 per cent of GDP, or even between $\frac{1}{4}$ and $\frac{1}{2}$ a per cent. But £1bn – large though it may seem – is by the end of the forecast a mere 0.03% of GDP.

That is well below the threshold at which we can have certainty about the precision of the forecasts. And yet often policies which might have significant benefits are dropped at the last minute to achieve the desired headroom. In its anxiety to portray numbers for political purposes, the Treasury appears to have mistaken specificity for accuracy and precision. It should instead decide what it wants its overall fiscal stance to be in broad terms rather than fixate on spuriously accurate figures. Restricting its opportunities to tinker with ever-expanding numbers of forecast rounds and broadening the indicators assessed by the OBR might just move us closer to a better paradigm.