

Signal versus noise

Speech by Ben Broadbent

Given at the London Business School

18 December 2023

Speech

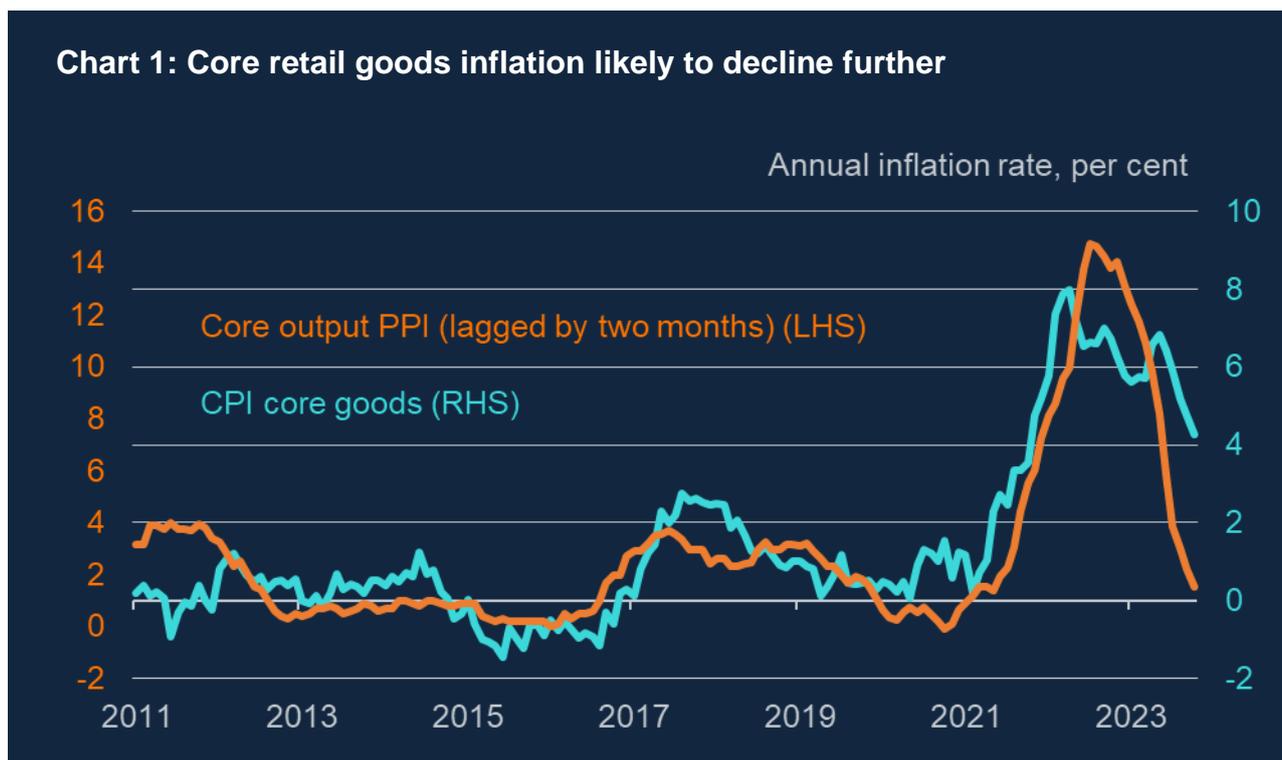
Introduction and summary

We are all aware that the future is unpredictable. When it comes to gauging the economy, however, it's not just the future that's uncertain: so is the present. For all the time and effort put into its forecasts the MPC also spends a great deal of it getting to understand the here and now.

This isn't straightforward. For one thing, it's not always possible to determine precisely what's causing what – to trace the economy we observe back to the underlying forces that are driving it. Is output growth being moved around by demand or supply? Ditto employment? To take an example of particular relevance right now, has strong wage growth been the result of exceptional tightness of the labour market, especially last year; or is it the “second-round effect” of very high spot inflation in late 2022 and earlier this year?

These things aren't mutually exclusive: almost certainly, both have played a part. But the balance of the two matters. As the direct effects of the pandemic and the war dissipate, wholesale prices of energy and other traded goods have been declining. This is now feeding through to inflation rates for retail goods prices and the aggregate CPI itself (Chart 1 plots core goods inflation against its wholesale counterpart; there have been similar trends in food and energy markets). As this happens, one might expect these second-round effects on wage growth and broader domestic inflation to weaken as well, quite independently of the stance of monetary policy. To the extent the tight labour market is the cause of strong domestic inflation, however, then the economy would need a longer period of below-trend growth – possibly with corresponding consequences for monetary policy – to bring it back into a more sustainable position. At any rate, the more general point is that it's not always easy to infer the deeper, unobserved causes of economic fluctuations from the directly observable information¹.

¹ Technically – and perfectly intuitively – a necessary condition for being able to infer (“identify”) these deeper causes is that we have at least as many observable variables. If (for example) you see only the quantity of output of some good but you know it can in principle be affected by two unobservable drivers (demand or supply shocks, say) then, in the absence of any other information, it obviously won't be possible to say which of them is behind any particular move in output. All you can do is to guess, based on their relative importance over the past, which is likely to have been the more significant (there's an example of this kind of “signal extraction” problem in a simple model below.) And even when that numerical condition is satisfied, and a more complete allocation of effect to cause is possible, estimation error in our models means we can't do that with complete precision. In the case of demand and supply shocks, for example, you can in principle work out which is behind any particular movement in quantities if you also get to observe what's happening with prices: increases in demand and supply both raise output but they have opposite effects (up and down respectively) on prices. However, to determine their respective contributions you still need to know

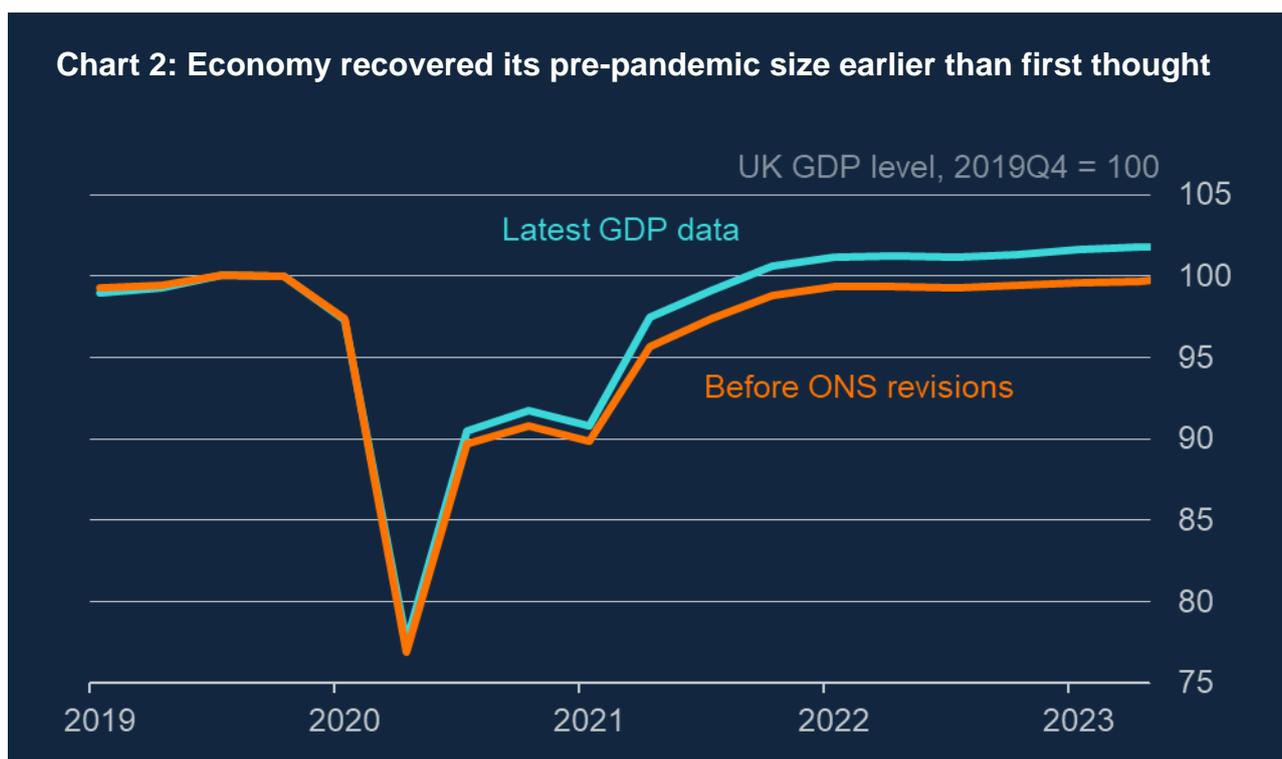


Sources: ONS and Bank calculations.

Core output PPI is for manufactured products excluding non-core items. Core output PPI annual inflation is advanced by two months, which maximises correlation with CPI core goods inflation.

Second, even what we do get to observe – GDP, employment, wages and the like – may not be perfectly measured. For some things (notably GDP) the relevant information comes in only over time and, as a result, the data are subject to revision. These changes can be sizeable. Recently, for example, estimated growth during 2020 and 2021 was revised up by almost two percentage points (Chart 2). As a result the economy is now thought to have reached its pre-pandemic size nearly two years earlier than was previously thought.

the slopes of the supply and demand curves and, unavoidably, there is a degree of uncertainty about exactly what those are.



Sources: ONS and Bank calculations.

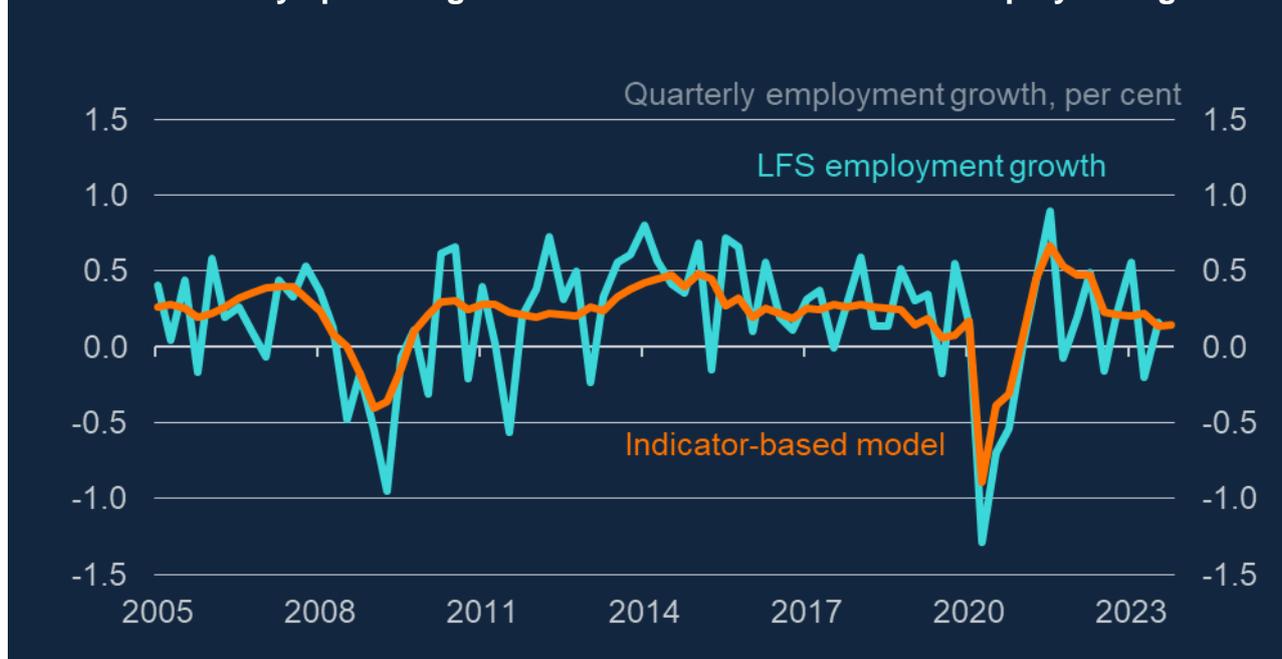
ONS 2023 Blue Book revisions. See Box C in the [November Monetary Policy Report](#) for more details.

Even in cases where revisions are less important, there's almost invariably a degree of measurement error in economic series. Recently, a steep decline in the response rate led the ONS temporarily to suspend publication of its Labour Force Survey, used to construct the official estimates of employment and unemployment². But, in the UK series, there was always some sampling error in official employment data, probably at least part of the explanation for the “zig-zag” pattern in quarterly growth rates (Chart 3³).

² Response rates had been declining before 2020 but, perhaps as a lasting effect of the pandemic, have done so faster since. Similar downward trends have been seen in many other countries (see, for example, [de Leeuw et al, 2018](#)). The ONS plans to re-start publication of the LFS in the New Year. For some time it has also been developing an improved, “Transformed Labour Force Survey” (TLFS). This should be ready for publication in the spring of next year. Details of the TLFS, and the ONS’s methods for constructing estimates of labour market quantities in the meantime, can be found [here](#) and [here](#).

³ If employment is over-measured one quarter then, on the reasonable assumption that the measurement error is uncorrelated from one period to the next, you’d expect on average to see weaker growth the following quarter. In symbols, if (the log of) employment is measured with some serially uncorrelated error ϵ , such that $\hat{l} = l + \epsilon$, where l is the true number, then its growth rate will have a moving-average error term: $\Delta \hat{l} = \Delta l + \Delta \epsilon$. In simple tests on the UK time series the presence of this term shows up with a high degree of statistical significance. Independently, and alongside the main LFS release, the ONS has always published estimates of sampling error for various series. The latest release, in September this year, had a figure of ± 0.5 percentage points for quarterly employment growth.

Chart 3: Surveys provide good real-time information about employment growth

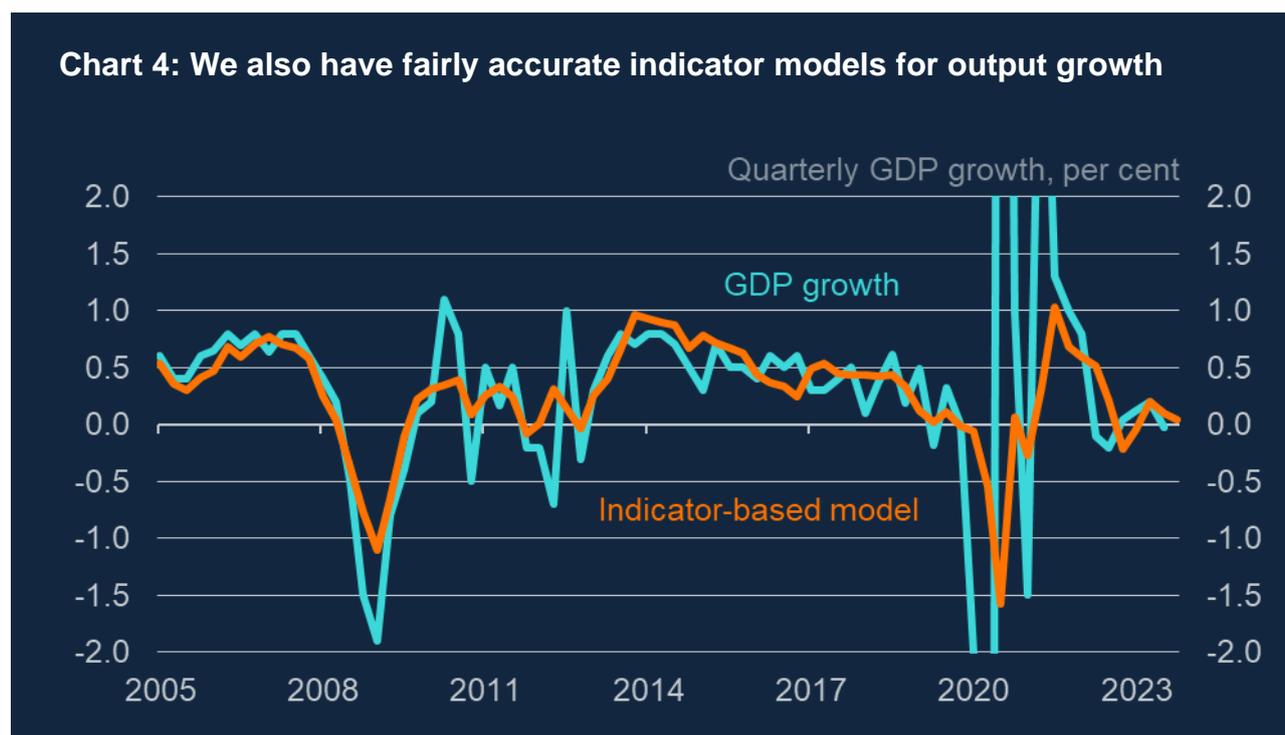


Sources: Bank of England, HMRC, KPMG/REC/S&P Global UK Report on Jobs, Lloyds Business Barometer, ONS, S&P Global/CIPS and Bank calculations.

From July 2023 LFS employment growth was replaced by ONS experimental Labour Market Statistics – see this [ONS Labour Market Overview](#) for more details. Latest data point is for October 2023.

Bank staff's indicator-based model of near-term employment growth uses mixed-data sampling (or 'MIDAS') techniques (see [Daniell and Moreira \(2023\)](#) for more detail). A range of indicators inform the model, including series from the Bank of England Agents, the Lloyds Business Barometer, ONS/HMRC PAYE payrolls, S&P Global/CIPS purchasing managers' index and the KPMG/REC UK Report on Jobs. Indicators are weighted together according to their relative forecast performance in the recent past.

This is why, in both cases, the MPC has always supplemented official estimates of GDP and employment growth with additional sources of information, usually from business surveys. The orange lines in Charts 3 and 4 are weighted averages of survey readings and other indicators, for employment and output respectively, designed to tell one something about their "underlying" rates of growth.



Sources: S&P Global/CIPS, CBI, BCC and Bank calculations.

GDP growth is the ONS's first estimate. Bank staff's indicator-based model of near-term GDP growth uses MIDAS techniques, like for employment growth. A range of indicators inform the model, including from the S&P Global/CIPS Purchasing Managers' Index, CBI Growth Indicator, and BCC Quarterly Economic Survey.

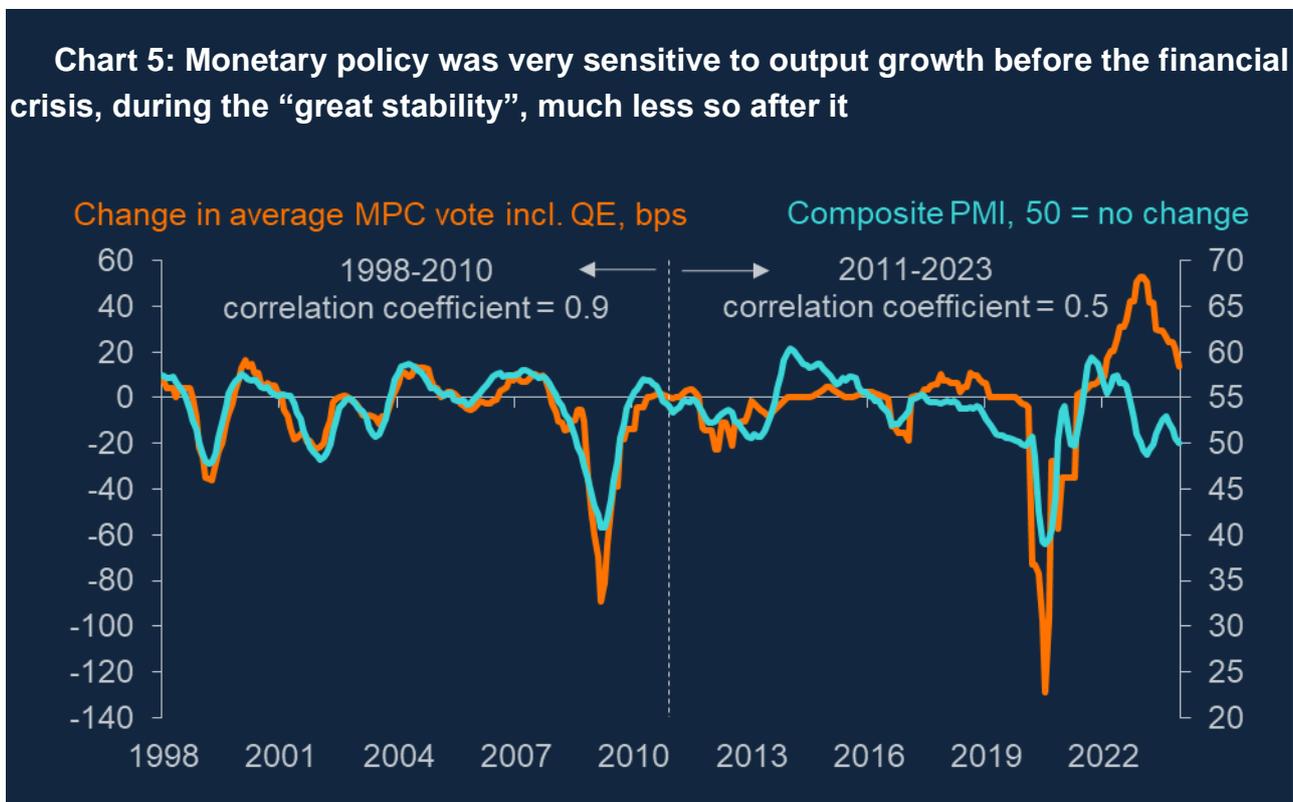
One shouldn't exaggerate the problem: in many instances these difficulties aren't that important. The indicator models in Charts 3 and 4 are actually fairly good. We have more sources of information than we once did and also better techniques for using it – including the “nowcasting” methods pioneered by Lucrezia Reichlin here at the LBS⁴.

Regarding GDP, errors in early estimates are in general smaller than they used to be back in the 1980s and 1990s. And long-dated revisions of the sort we saw in this year's Blue Book, relating to economic growth well over a year ago, won't in any case have much bearing on monetary policy as any effects the additional output might have had on employment and subsequently inflation would already have come through. (As such, they're more likely to affect one's view of the level of potential supply in the economy, rather than the output gap (or, therefore, future inflation)).

As for employment, there was a time when the MPC seemed to put very little weight on news from the labour market, however well measured the data. Indeed in the pre-financial-crisis period it was hard to find anything that mattered for interest-rate decisions, empirically speaking at least, other than output growth. Chart 5, which I've used several times in the past, is one striking way to get the point across. It plots a survey measure of

⁴ See, for example, [Giannone, Reichlin and Small \(2008\)](#). The particular models used by the Bank are described in [Anesti et al \(2017\)](#) and [Daniell and Moreira \(2023\)](#).

GDP growth against the average vote for changes in Bank Rate, in basis points, on the MPC⁵. The pre-2010 correlation is extremely tight and, once one's allowed for it, nothing else – whether unemployment, wage growth or inflation itself – offers any further statistical help in explaining how policy rates behaved.



Sources: Bank of England, S&P Global and Bank calculations.

Six-month rolling averages.

During that pre-crisis period, when policymakers were confident about the stability of underlying productivity growth, and of the supply side more generally, this made sense. An acceleration in demand would lead to lower unemployment, albeit with some delay; as long as the NAIRU is also stable, lower unemployment would necessarily, in time, indicate greater inflationary pressure, including faster wage growth; and with inflation subject to few other disturbances, independently of demand, why wait until the end of this chain to act? You may as well tighten policy as soon as you see the economy strengthen and, symmetrically, loosen when output growth declines. In this happy environment, sometimes known as the “divine coincidence”, the policy that best stabilises economic growth is also the policy that offers the best chance of controlling inflation.

The financial crisis marked the end of this “great stability” and with it this simple feedback rule for monetary policy. Those days are long behind us. As Chart 5 reveals, policy

⁵ Using a crude multiplier to turn one into the other, the red line includes the Bank-Rate equivalent of votes for QE.

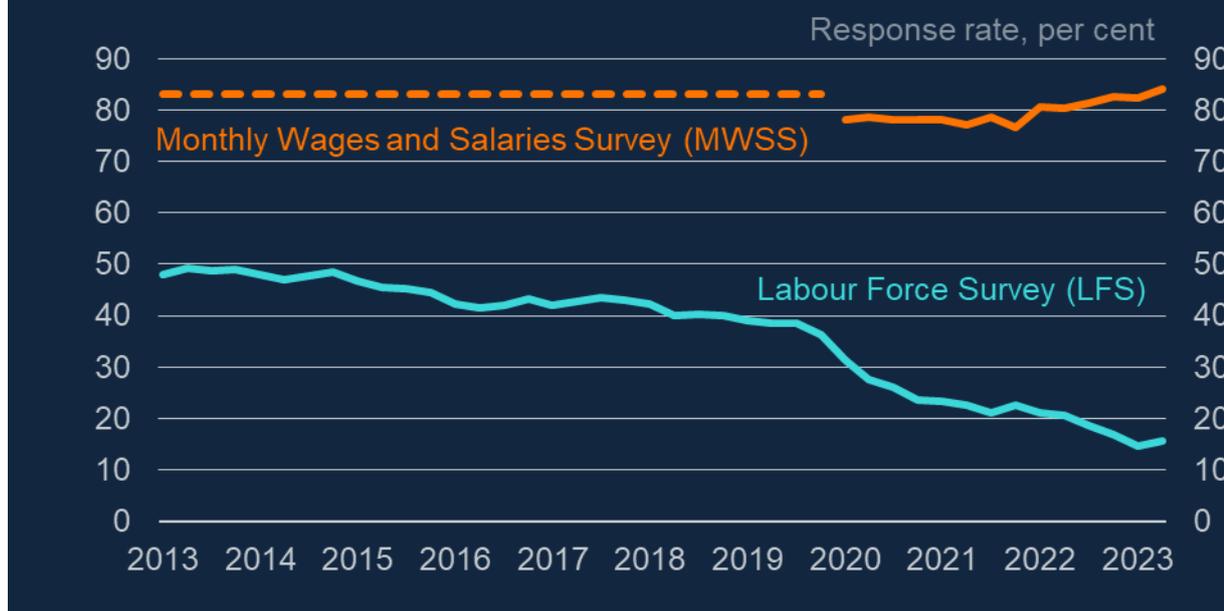
decisions have since been much less well correlated with GDP growth. Instead – or least in addition – policy then started responding to more direct measures of the degree of slack in the economy, in particular the rate of unemployment. Latterly, the MPC has also had to put weight on nominal indicators such as wage growth and other late-cycle measures of domestic inflation.

Shortly I'll try and explain why this is. This will basically be a restatement of a point I made in a talk in the summer ([Broadbent, 2023](#)) but, to flesh it about a bit, I'll chuck in the results of a simple model. This is designed to get across how policy should behave when there's incomplete information about the various shocks hitting the economy.

I'll then take a closer look at the unemployment and wage data. If the absence of the Labour Force Survey (LFS) employment estimates isn't so serious, the temporary absence of the series for unemployment is felt a little more keenly. It matters more in models of wage growth. And, during the temporary suspension of the LFS, the ONS is in the meantime using a proxy – the claimant count – that may not be quite as accurate as those we have for employment.

As for wages, one should say straightaway that the source of information for the official average weekly earnings (AWE) series, a survey of businesses called the Monthly Wages and Salaries Survey (MWSS), has not suffered from the same declining response rates as the LFS (Chart 6). That said, it is a rather volatile series, particularly from quarter to quarter, and at times this year there has been unusually high disparity between the AWE and other measures of wage growth that the MPC routinely looks at.

Chart 6: The survey used for official estimates of wage growth hasn't suffered from the same declining response rates as the LFS



Sources: ONS and Bank calculations. As shown in [Haskel, 2023](#).

Data show calendar quarter averages. Before 2020 Q2, the dashed line shows the MWSS target and pre-Covid norm response rate of 83%, as reported by ONS. Latest data point July-September 2023.

Why unemployment and wage growth matter

As we saw a moment ago, policy in the pre-crisis “great stability” responded very sensitively – and pretty much uniquely – to economic growth. This can be justified if the policymaker is confident about the stability of the supply side of the economy.

But if the financial crisis marked the end of a long period of healthy productivity growth, it also made the supply side more uncertain. And this in turn necessitated a change of approach for monetary policy.

When underlying productivity growth can vary, and swings in output might be associated as much with supply as demand, they may no longer carry the same implications for the degree of slack in the economy. In that case it makes sense to put at least some weight on more direct indicators of spare capacity, including unemployment, even if one loses some time in doing so. That time is lost because of the delay between shifts in demand and their consequences for unemployment (in UK data the figure seems to be around 4-6 months). Following a genuine demand-driven rise in output, waiting to see what happens in the labour market would prevent you from responding to it as promptly as you otherwise would. On the other hand it might also save you from mistakenly tightening policy if it turns out to have been driven (or at least accompanied) by better supply. As long there's a

reasonable degree of stability in the NAIRU – the rate of unemployment consistent with stable wage growth and domestic inflation – then changes in unemployment can give you a cleaner (if later) read than GDP growth alone on how spare capacity is developing. This, at least, is how I understood the MPC’s conditional guidance back in 2013 (which singled out the rate of unemployment) and, more generally, the weight the Committee began to put on the labour market following the financial crisis ([Broadbent, 2014](#)).

However, if even the NAIRU can shift around – or, more generally, wages and domestic inflation are subject to important influences other than spare capacity – then GDP and unemployment may still be insufficient. You also need to consider the behaviour of these nominal variables more directly. As I sought to explain in a talk this past summer, this is one way to understand why the MPC has had to pay more attention to wage growth and services inflation than in the past – more, in fact, than to economic growth.

Judging by the huge expansion in the number of vacancies in 2022, and steep declines too in surveys of labour availability, the labour market seemed to be significantly tighter than the rate of unemployment alone would have suggested (Chart 7).

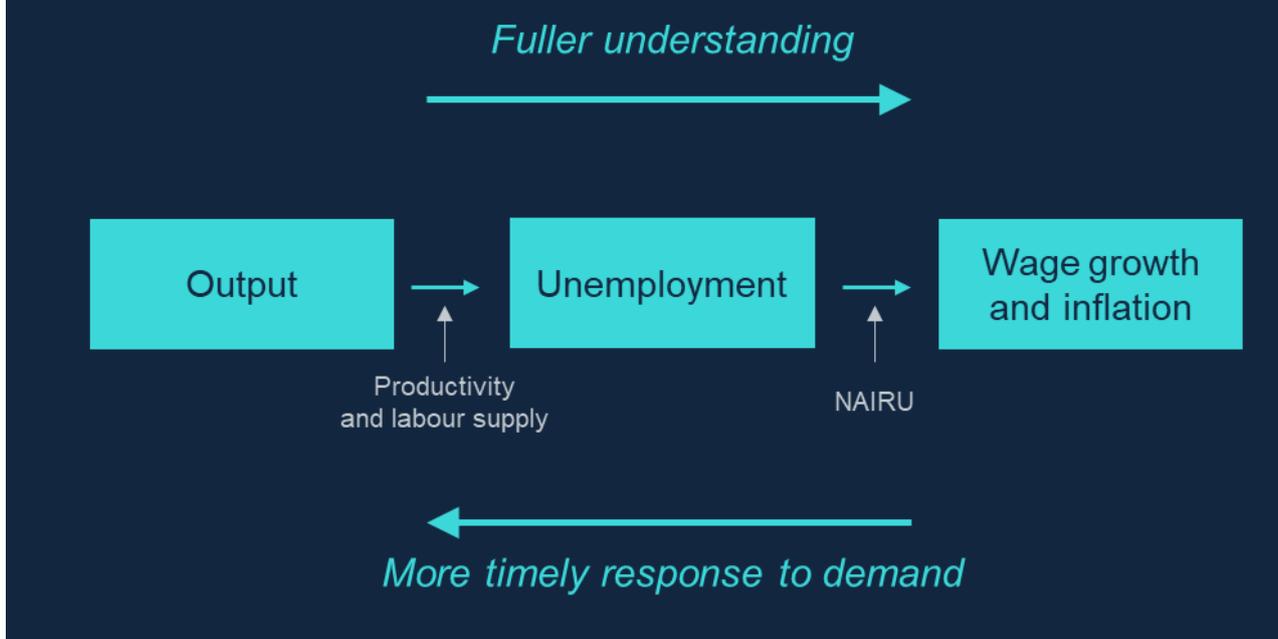


Sources: Bank of England, ONS and Bank calculations.

This indicated that that the NAIRU had risen, perhaps because, amidst all the shifts created by the pandemic, the labour market was having a hard time reabsorbing the

significant numbers coming back to work and reallocating them in line with demand⁶. The very pronounced second-round effects of higher import prices on domestic prices and wages can be thought of in a similar way – they too can boost inflation even for a given rate of unemployment. Taken together these developments suggest that, on top of output growth and unemployment, policy should put some weight on these nominal variables, even if that means waiting that much longer to respond to (what subsequently turns out to have been) true shocks to demand. Trading off timeliness for a better understanding of what’s going on in the economy, uncertainty about the supply side pushes one further to the right of the schematic picture in Chart 8.

Chart 8: Lags in the inflation process mean that, when the policymaker is uncertain about supply, there’s a trade-off between a better understanding of the economy and the timeliness of the policy response

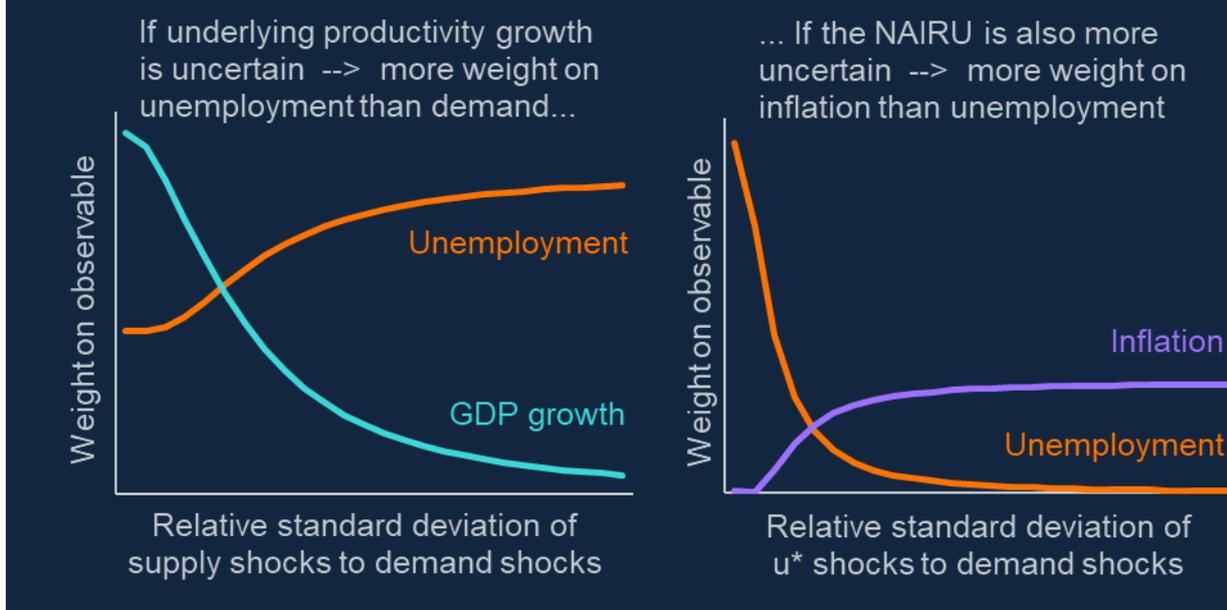


The lag between movements in spare capacity and inflation is longer and the price one pays for this extra information, in time, is therefore higher. There are clearly risks in moving only when you see “the whites of inflation’s eyes”. But it may still be a price worth paying. The two graphs below (Chart 9), derived from a simple model of inflation in which

⁶ It was also the case in many countries, including the UK, that the pandemic resulted in a significant decline in labour market participation, persuading some people to retire earlier than they might have done. It’s not obvious this matters directly for inflationary pressure, at least once you’ve taken into account any impact it might have had on the rate of unemployment: in most simple empirical models of wage growth, it’s unemployment (relative to the NAIRU) that is the relevant measure of slack in the labour market. However, it’s possible the early retirees were more concentrated in some parts of the economy than others, in which case the drop in participation might have added to mismatch in the labour market, contributing to the rise in the NAIRU.

there's incomplete information about the shocks hitting the economy⁷, help to cement the point. In this representation inflation can be moved around by four underlying drivers – demand, productivity, firms' mark-ups and movements in the NAIRU – but the policymaker has only three observable variables from which to deduce them (output, unemployment and inflation itself). Importantly, he or she is assumed to understand how variable these underlying disturbances are (estimates that can be built up over time). But there is nonetheless what is sometimes called a “signal extraction” problem, and one can estimate only imprecisely, in real time, what is driving what. And, in doing so, the policymaker will lean on a greater range of information.

Chart 9: More weight put on unemployment when there's uncertainty about underlying productivity growth; uncertainty about the NAIRU means wage and price inflation matter too



Sources: ONS and Bank calculations.

First panel plots the contribution of activity and unemployment data to the estimated variance of the output gap, where along the x-axis we increase the standard deviation of the supply shock relative to the others.

Second panel charts the contributions of unemployment and inflation to the estimated variance of the output gap, where along the x-axis we increase the standard deviation of the shock to u^* .

The graphs plot two pairwise comparisons of the relative importance of the observable series for estimates of the output gap, depending on the variance of (and therefore uncertainty about) the underlying shocks. As uncertainty about underlying productivity growth rises, the policymaker puts increasing weight on developments in unemployment rather than GDP growth. When the NAIRU is more uncertain, inflation becomes more

⁷ The model is set out in more detail in a short Annex to this speech.

significant. As that happens, the variability of inflation itself rises, because policy has a less precise understanding (at least in real time) of what's driving the economy and there's a greater chance of a misdiagnosis⁸.

This is a pretty simple representation but it does, I think, capture something about the environment we've been in for the past two to three years.

Greater uncertainty about movements in unemployment

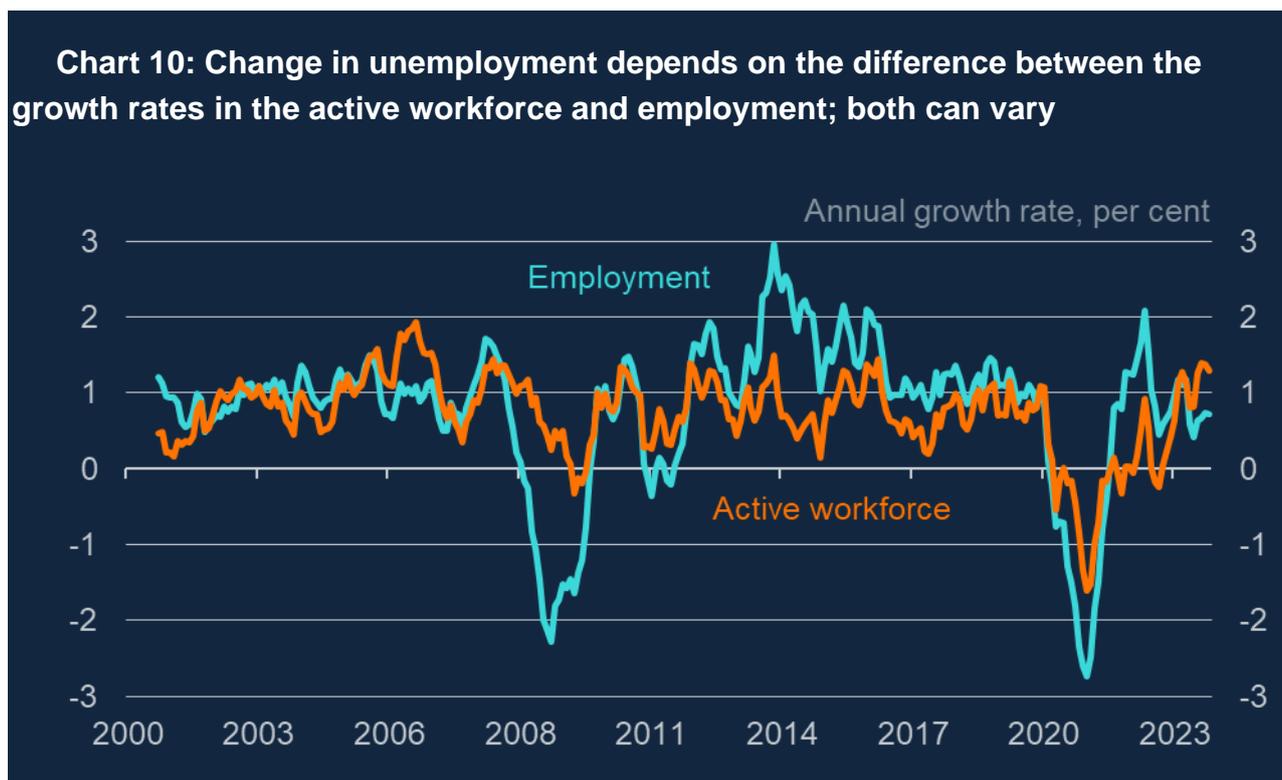
Having tried to get across why indicators like unemployment and wage growth have been important for policy, I now want to take a closer look at the data themselves.

I mentioned earlier that business surveys can give one a reasonable read on what's happening to employment. To that extent the temporary suspension of the LFS isn't that costly.

I don't think the same is true of unemployment. If the active workforce were reasonably stable it wouldn't be too much of a problem. In that case one could infer changes in unemployment directly from those in employment.⁹ In practice, however, estimates of labour market activity move around quite a bit, even over relatively short periods of time (Chart 10). The most recent Labour Force Survey, in quarterly data, was for 2023Q2. Over the previous three quarters, from its trough in the third quarter of 2022, the rate of unemployment was estimated to have risen by a little over half a percentage point. That happened not because employment declined – it actually grew reasonably strongly – but because of a recovery in participation. According to the LFS, significant numbers of people decided to return to work and not all of them were able to find jobs straightaway. The more general point is this: we care more about unemployment than employment and movements in the second aren't necessarily enough to tell us about the first.

⁸ In a speech she's due to give tomorrow, my colleague Sarah Breeden presents some simulations along these lines, exploring what happens when the policymaker is uncertain about the degree of persistence in inflationary disturbances. If policy is therefore set assuming some sort of average across the various possibilities, and the degree of persistence turns out to be towards the upper end of that prior distribution, inflation will obviously end up overshooting the target. Similarly, if the shock turns out to be more transient, inflation will end up being too low.

⁹ If u is the rate of unemployment, g_A the growth in the active workforce and g_E that of employment, then the change in unemployment is $\Delta u = (1 - u)(g_A - g_E)$. This differs from the difference between the two growth rates only by a factor $1 - u$, which is in general is pretty close to one.

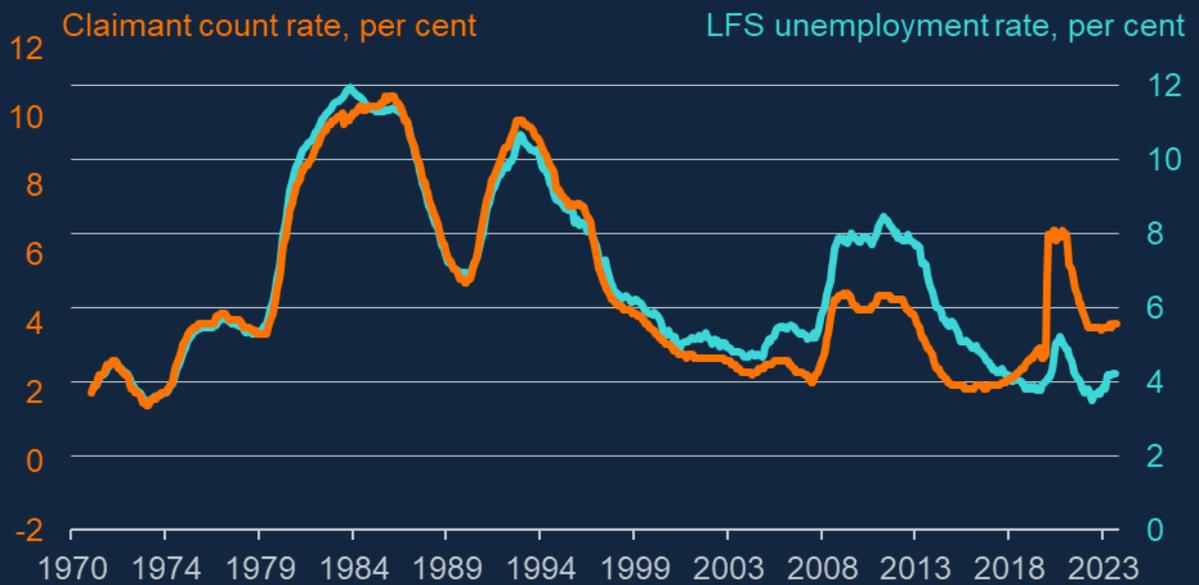


Sources: ONS and Bank calculations.

So, during this temporary suspension of the LFS, the ONS is instead basing its estimates of unemployment on developments in the claimant count – the numbers of people receiving unemployment-related benefit.

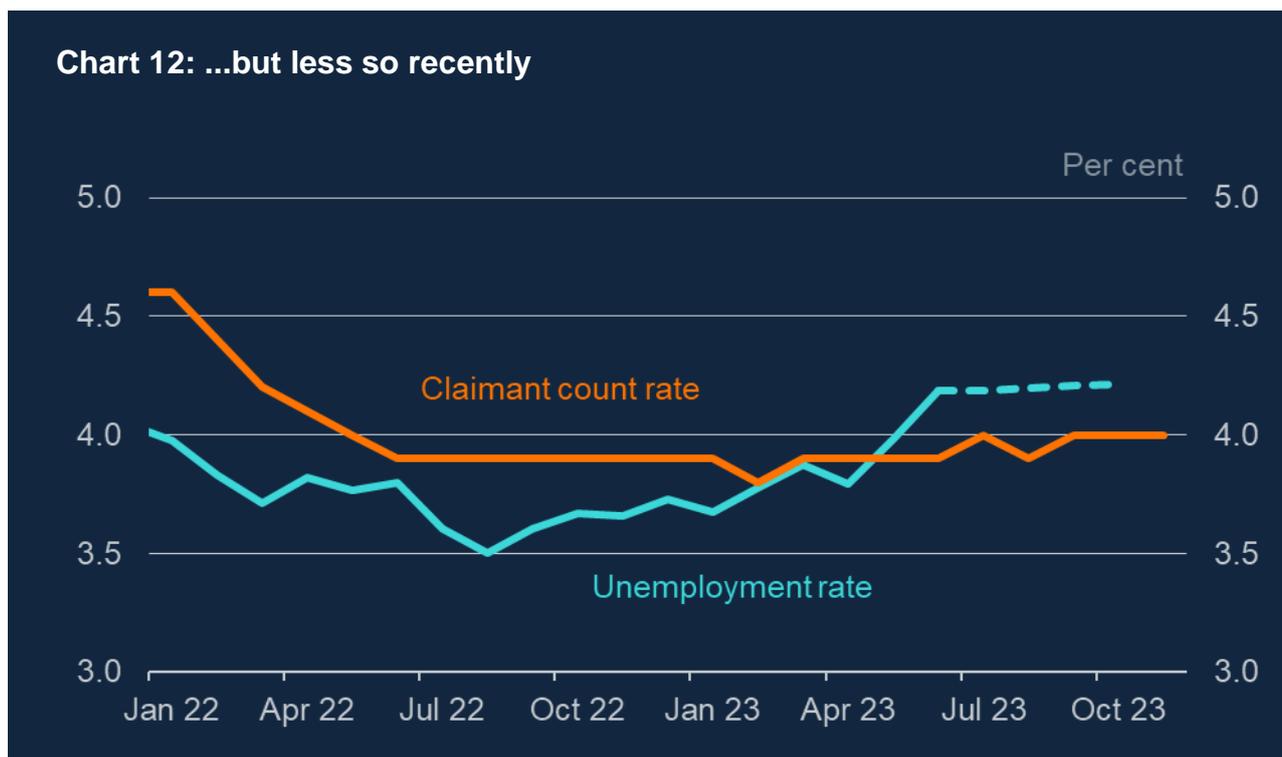
However – and even if there's nothing better – it's not clear this is a perfect proxy. People can and often do look for work without claiming benefit. And although the two measures are reasonably well correlated in the past, and over longer periods of time, that's less true from quarter to quarter, particularly in the past couple of years (Charts 11 and 12). The rise in LFS unemployment through the first half of 2023 wasn't matched by any sort of increase in the claimant count.

Chart 11: Unemployment and claimant count closely related in the past...



Sources: ONS and Bank calculations.

The correlation between unemployment rate and claimant count rate fell with the roll-out of Universal Credit, which increased the number of people eligible to claim benefits. As part of the UK government's response to COVID-19, more people became eligible for unemployment-related benefits, although still employed. As a result, changes in claimant count were not due wholly to changes in the number of unemployed people.



Sources: ONS and Bank calculations.

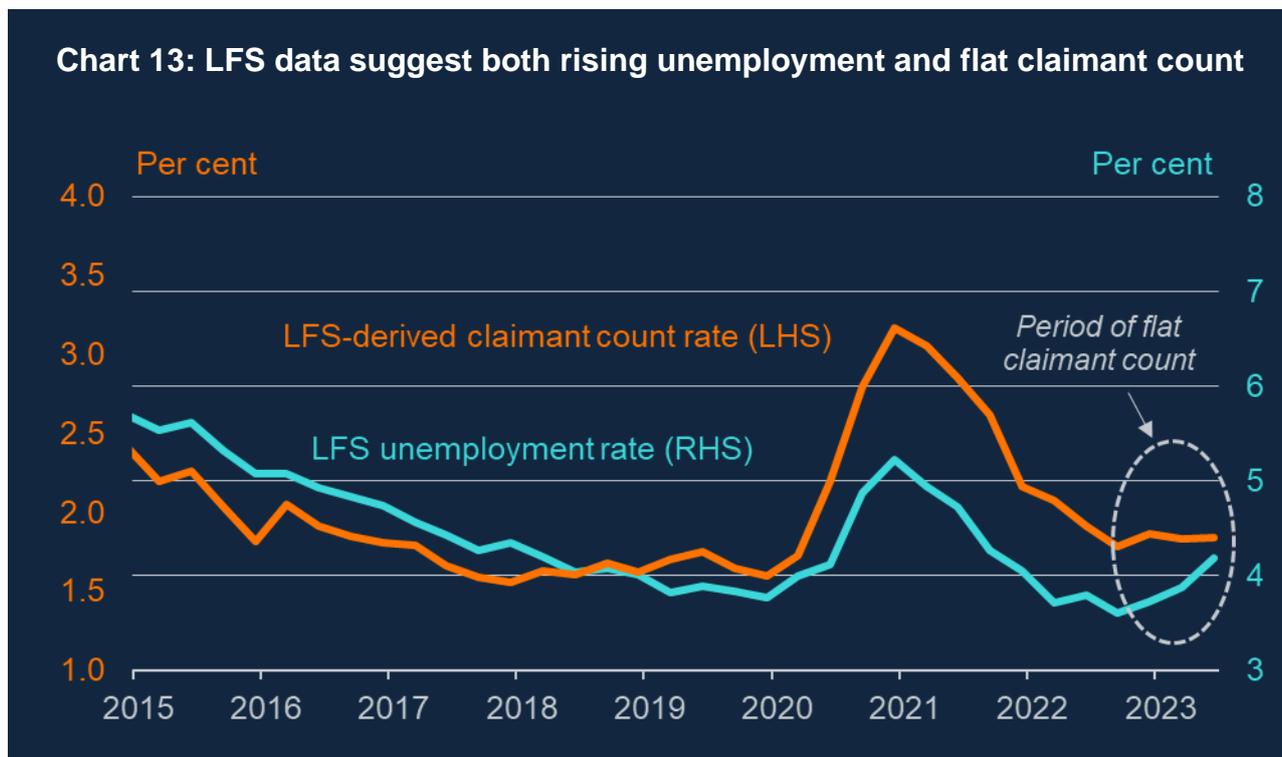
Unemployment rate published as an experimental statistics (shown as a dotted line) starting from July – see this [ONS Labour Market Overview](#) for more details. Unemployment data available to October, claimant count data to November.

It's possible that the problem exists not in the benefit numbers but the LFS. After all, if there's measurement error in the level of employment, presumably the same must be true of unemployment. So maybe the benefit numbers are telling us that the rise in LFS unemployment, even as moderate as it's been, has been overstated.

But for several reasons I suspect it's the other way around – that the claimant count understates the rise in unemployment over that period. First, for what it's worth, LFS unemployment is considerably less variable over short periods of time than employment growth. This suggests that its measurement error may be less acute.

Second, the picture of rising unemployment but flat claimant numbers is one that's actually confirmed within the LFS itself. One of the more detailed questions in the survey is whether or not the respondent is receiving unemployment-related benefit. And that series too has been flat (Chart 13). Third, there are good reasons, in periods of rising participation, why one might expect the claimant count to understate any rise in actual unemployment. The detailed numbers in the LFS reveal that people looking for work are much more likely to claim benefit if they'd had a job immediately beforehand. Those moving from inactivity – people coming back into the labour market after taking time off, for example, or students leaving full-time education – are much less likely to do so (Chart 14), even if they're unable immediately to find work and therefore count as unemployed. This

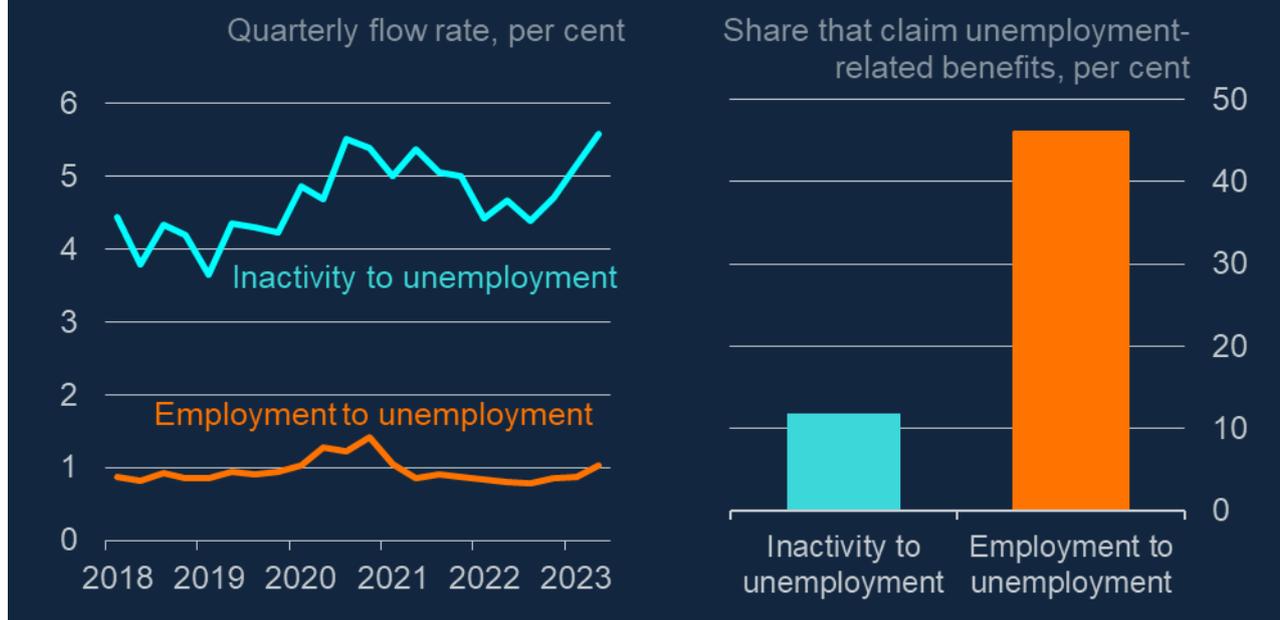
makes intuitive sense. And it means that one wouldn't expect any material rise in the claimant count unless or until there were significant job losses.



Sources: ONS and Bank calculations.

In the LFS unemployment-related benefits are defined as claims for Job Seekers Allowance or Universal Credit for the reason of being 'unemployed and searching for work'. This is a slightly narrower definition than for the headline claimant count measure. Latest data point is Q2 2023.

Chart 14: People who've lost jobs much more likely to claim benefit than those (re-)entering the labour market from inactivity



Sources: ONS and Bank calculations.

First panel plots quarterly labour market flows data for people aged 16-64. Flows relate to activity status three months prior. Right panel plots share that report claiming unemployment-related benefits. Activity status relates to 12 months prior. Latest data point is Q2 2023.

I should say that there's very little sign of that. The redundancy rate ticked up a little during the first half of the year. So-called "HR1 notifications" – advance notice of redundancies that larger employers are obliged to give to the government's Insolvency Service – provide somewhat timelier information and those too have edged up. But they're still very low by historical standards.

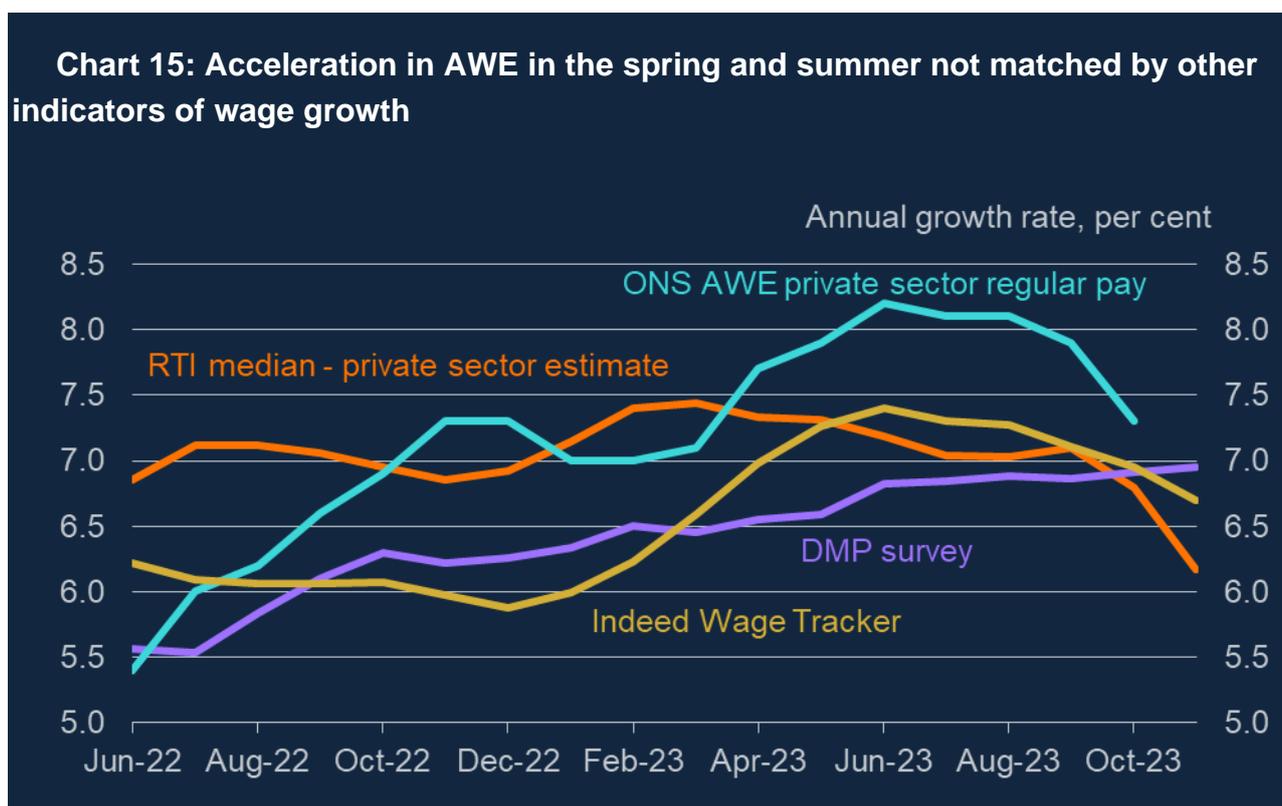
More generally, I do not mean to imply that the true rate of unemployment is definitely still rising. Our indicator model suggests employment is broadly flat so, at least in the near term, any rise in unemployment would arise only because of continuing growth in the active workforce – and the participation rate is simply too volatile to count on that in the near term.

The point I want to make is more that, during the temporary suspension of the LFS, the loss of the unemployment series matters a little more than that for employment. It matters more in basic models of wage growth (which is why it's one of the variables stressed by the MPC). And proxies for the change in unemployment – whether just flipping round the figure for employment growth or seeing what's happened to the claimant count – do not, over short periods of time, perform that well.

Wage growth: how high did it go and is it now declining?

As we saw earlier, the survey used to construct the official wage series has none of the response-rate problems that have afflicted the LFS (Chart 6). A high and stable proportion of businesses asked to do so complete the MWSS. It covers only those businesses that pay VAT or PAYE (about 120,000 of the 5-million-plus registered firms in the UK); of those “only” 9,000 complete the survey. But the skew – in both the eligible and surveyed populations – is very much towards the larger businesses. So the AWE actually captures 14 million employees, almost half of their total number.

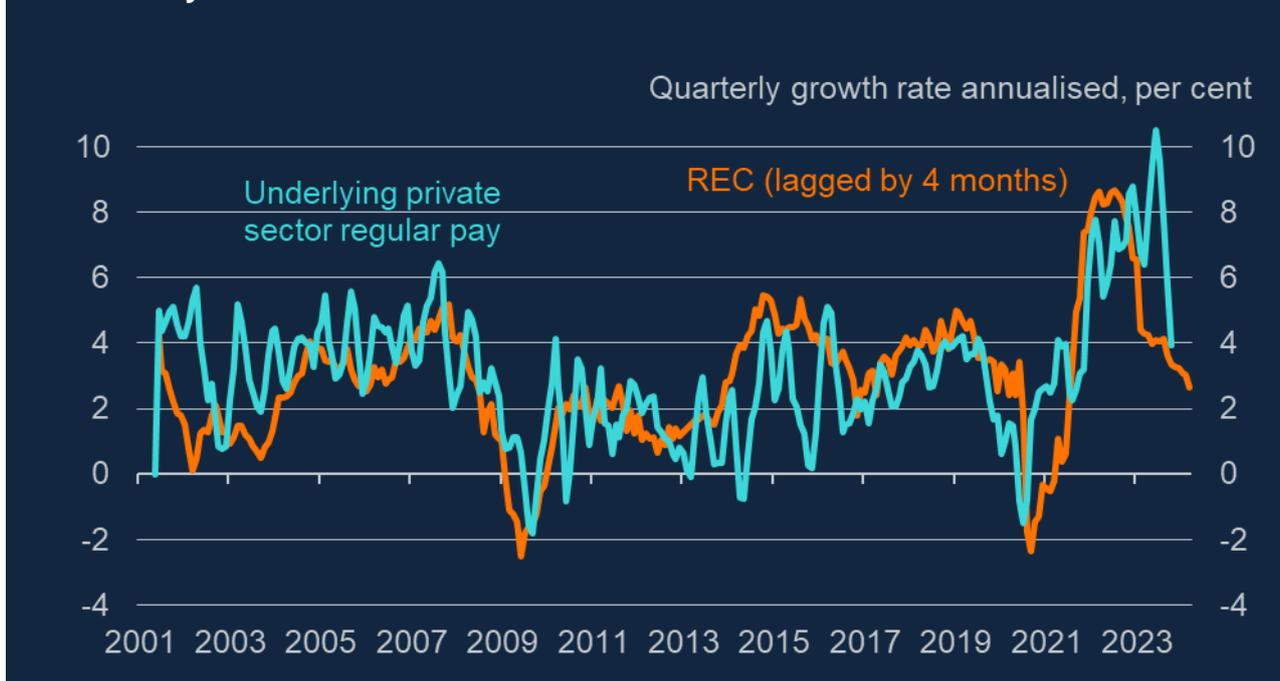
All that said, the official wage data are nonetheless subject to revision, thanks to late returns; even though it's adjusted for outliers (the ONS tops and tails the raw numbers from the MWSS) AWE growth can be pretty volatile, particularly over short periods of time; and, for much of this year, the headline figure has been higher than many other indicators of wage growth (Charts 15 and 16).



Sources: DMP Survey, Indeed Hiring Lab, ONS and Bank calculations.

All measures are three-month average over three-month average a year earlier. RTI ‘private sector’ pay is a Bank staff estimate that strips out sectors with a high share of public workers, such as public administration and defence, social security, education, health and social work.

Chart 16: The REC has been a good leading indicator of AWE growth for much of its history but not in 2023



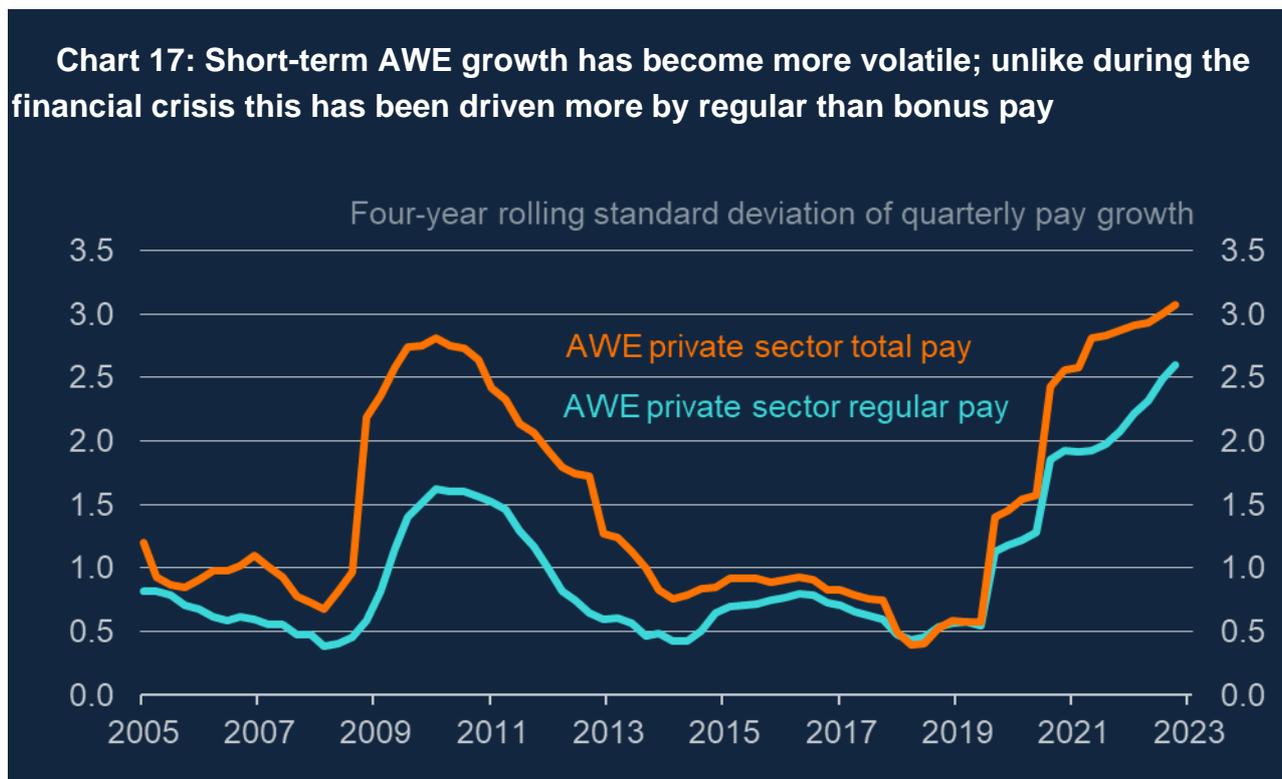
Sources: ONS, REC/KPMG/S&P Global Report on Jobs and Bank calculations.

Underlying private sector regular pay growth is Bank staff's estimate of underlying pay growth between January 2020 and November 2022 and ONS AWE private sector regular pay growth otherwise. REC measure is an index of permanent staff salaries, mean-variance adjusted to quarterly ONS private sector regular pay growth over March 2001–19. REC is advanced by four months, which maximises correlation with AWE quarterly private sector regular pay growth.

It may be that we're now seeing a degree of re-convergence – the official data have softened a bit in the last couple of months. But the slightly muddy picture of the recent past, coupled with the general volatility of the data, means the MPC would probably want to see more evidence, across several indicators, before concluding things are on a clear downward trend.

The recent volatility of the AWE is apparent enough in Chart 15. It can be seen more directly in Chart 17, which plots the standard deviations of quarterly private-sector AWE growth over rolling four-year periods (so that the very last data point covers the period since the onset of the pandemic). The sharp pick-up in wage growth has been accompanied by a similar rise in its volatility. It's also clear in Chart 18. This plots the output of what's called a "dynamic factor model", a statistical filter that uses the sectoral information within the AWE to decompose headline wage growth into "trend" and

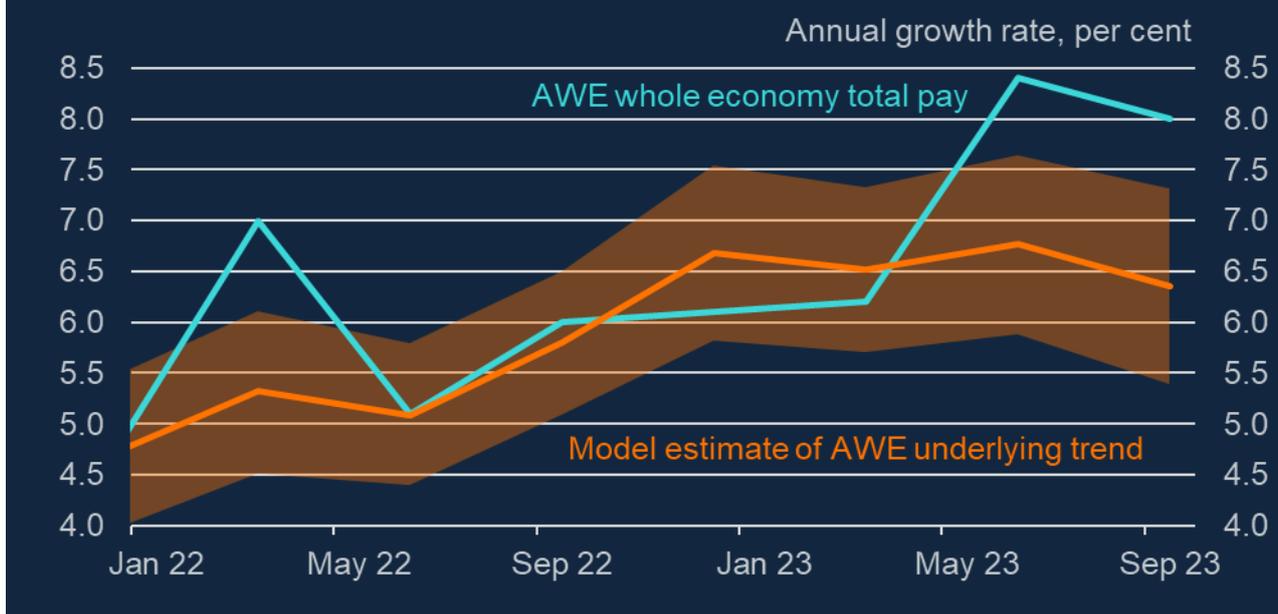
“transitory” components¹⁰. Quite a bit of the marked acceleration in AWE growth over the spring and summer is interpreted, at least by this model, to have been transitory.



Sources: ONS and Bank calculations.

¹⁰ This model was estimated by economists at the Bank and is based on the approach of [Stock and Watson \(2016\)](#). Stock and Watson designed it originally as a measure of trend price inflation, using 17 sectors of the US PCE inflation index.

Chart 18: An estimate of trend AWE growth derived from sectoral information has been lower than the headline rate in recent months



Sources: ONS and Bank calculations.

Model of trend wage growth based on the approach of [Stock and Watson \(2016\)](#). Shaded area shows 68% confidence intervals around underlying trend estimate.

Unsurprisingly, in this environment, the MPC has been keen to supplement these official data with other indicators of wage growth. None of the three in Chart 15 has a long history¹¹; the DMP and the two surveys of recruitment agents (the Indeed wage tracker and the REC index in Chart 16) all have much smaller coverage than the official MWSS. So they should certainly be treated with a degree of caution.

Taken together, however, it's noticeable that, like the "trend" estimate in Chart 18, none exhibited quite the same strength as (or pick-up in) AWE growth in mid-2023. And I do think it's worth paying attention to two in particular, the REC index and the tax-derived "Real Time Information" (RTI) pay series.

As I said, the coverage of the REC survey must be significantly smaller than that of the AWE (most people don't find jobs via recruitment agents). But it's also much longer-lived than the other indicators (it started in 1997). And for nearly all that period it has had a pretty good – and, indeed, a leading – statistical relationship with short-term changes in the official wage index. Movements in the REC have tended to foreshadow those in quarterly AWE growth around four months later, with a reasonable degree of reliability.

¹¹ The tax-based RTI measure exists since 2014. The Indeed Wage tracker began only in 2019 and the DMP in May 2022.

This was clearly not the case earlier this year: the REC index peaked over a year ago and then fell back steeply over the following six months. On past form one might then have expected a decline in private-sector AWE growth starting early in the New Year. In fact, between January and July, the index accelerated significantly, rising at an annualised rate of over 9% over that six-month period.

It's not clear why the REC proved so misleading. One possibility is that the "second-round" effects of high inflation were skewed towards existing rather than new employees (coming from a survey of recruitment agents the REC picks up only the latter). We know that firms were accepting significant increases in wages early this year in order to help compensate employees for the steep rises in the cost of living. If these were aimed more at retaining than attracting people, and therefore paid disproportionately to people already at the firm, this would help to explain the divergence in Chart 16. It would also suggest that, as these effects fade, the relationship between the two should re-establish itself (certainly it's hard to think of something that would permanently have ruptured the link).

The second is the RTI series derived from tax data. This is much younger than the REC survey (it started in 2014). On the other hand, its sample size is even bigger than the MWSS – more than twice as big, in fact. It covers everybody subject to PAYE income tax (around 30 million people¹²). Helpfully, it also has a wealth of detail. One can distinguish between changes in the median and the mean (for various reasons economists at the Bank tend to favour the former¹³). It also offers a sectoral breakdown, allowing us to construct a reasonably good proxy for private-sector pay alone¹⁴. One drawback is that, unlike the AWE, it's not possible to strip out bonus payments (the MPC generally cares more about regular, non-bonus pay growth¹⁵).

Given these huge samples, for both series, and what must therefore be a big overlap between the two populations, it's perhaps surprising that the AWE and RTI series aren't routinely closer to each other. At least in 2023, some of the gap in Chart 15 may relate to this question of bonuses. According to the AWE, bonus growth was relatively subdued this year. So growth of the aggregate – with which the RTI series is probably more fairly

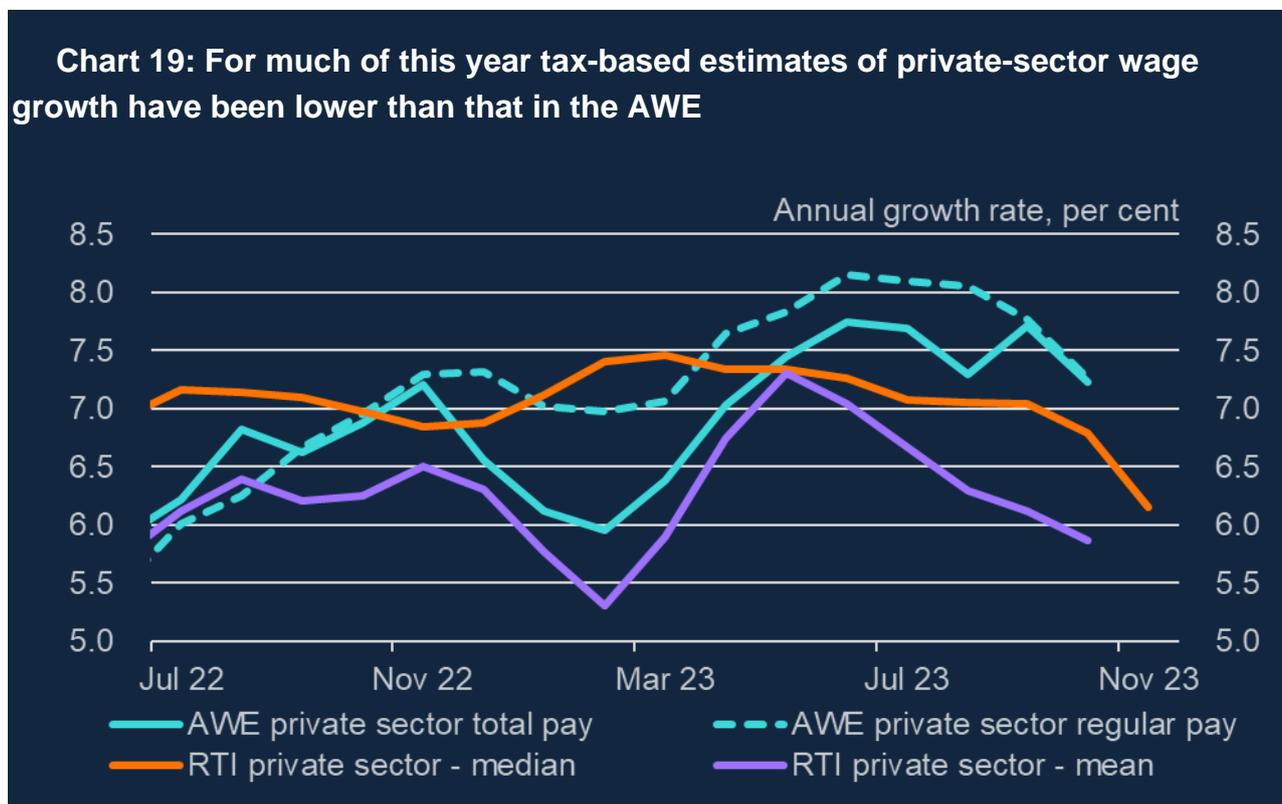
¹² This is greater than the number of employees (28.5 million). The reason is that the RTI statistics have slightly larger coverage (they include Northern Ireland, HM Armed Forces and government-supported trainees paid via PAYE, pay-rolled redundancy payments and signing-on fees, all of which are excluded from AWE). In addition, some people running their own businesses, and therefore classified as self-employed in the LFS, will still pay themselves salaries and therefore PAYE income tax as well.

¹³ The AWE measures mean pay. But, to construct it, the ONS lops off the outliers in the responses, pushing the measure closer to a median. The RTI statistics are experimental and may not be similarly adjusted.

¹⁴ The series for "private-sector" RTI pay growth in these graphs simply strips out industrial sectors with a wholly or near-wholly public-sector presence: public administration and defence, social security, education, health and social work.

¹⁵ On the whole the MPC is more interested in the regular (non-bonus) component of wages. This is partly because bonuses are more volatile. More importantly, it's not clear they really count as a marginal cost for the firm. Bonuses get paid out of profits already earned and, unlike regular pay, do not need to be pledged in advance as a condition of employment.

compared – was slightly weaker than that of regular pay, and therefore closer to the RTI figure (Chart 19). But there’s still a gap, although that too has narrowed a bit in the past couple of releases. In the three months to October, private-sector pay including bonuses was 7.2% higher than in the same period of 2022. This compares with 6.8% for the median of the RTI series, a figure that in the three months to November fell to 6.2%.



Sources: ONS and Bank calculations.

RTI ‘private sector’ measures are Bank staff estimates – see footnote to Chart 15.

I should be clear that the MPC still puts most weight on the AWE series. And it’s evident from all of them that pay growth – and with it other indicators of domestic inflation like the services component of the CPI – has been significantly faster than the sorts of rates that, over time, one might think of as consistent with the 2% inflation target.

But, taking all this together, there is a question about whether rapid pick-up in wage growth in the spring and summer was quite as marked as in the official data. AWE growth has been particularly volatile in the past year or so and other indicators of wage growth haven’t been quite as strong.

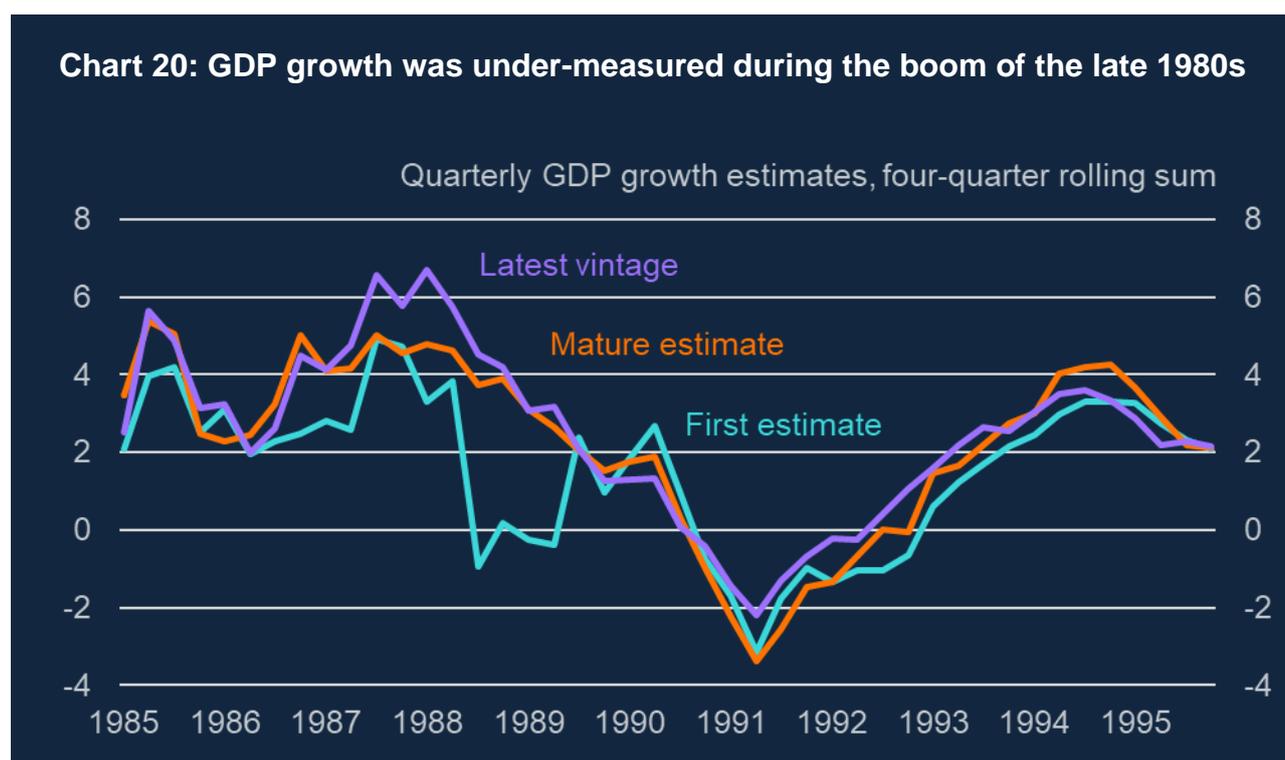
By the same token, of course, the noisiness of these measures, and the discrepancies between them must also inject a degree of caution about the more recent decline. One would want to see further evidence, across several indicators, before concluding things are on a clear downward trend.

Conclusion

In a speech in the autumn of 2007 my predecessor Charlie Bean said policymaking had sometimes been likened to “driving along a winding road looking only in the rear-view mirror” ([Bean, 2007](#)). If this was a warning about the unpredictability of the future then the events of the following year certainly justified it.

But it’s not just the future that’s uncertain – so is our assessment of the present and recent past. Charlie went on: “I wish it were that easy. In practice, the rear window is also a little misted up. Not only do we not know where we are going, but we have only an imperfect idea of where we have been.”

At times this uncertainty can be very costly. Early estimates of GDP growth during the late-80s boom significantly under-estimated the strength of the economy (Chart 20). This may have affected the setting of monetary policy and contributed to the scale of the subsequent inflation.



Sources: ONS and Bank calculations.

Mature estimate is three years after the first estimate.

Early GDP estimates have improved considerably since then. Because we also have more sources of information about economic activity, and better statistical methods for using them, inaccuracy in the measurement of current output growth is much less of a problem than it once was.

At the same time, uncertainty about the behaviour of the supply side of the economy has probably reduced the relative importance of these output data anyway. Instead (or at least in addition) monetary policy has had to put more weight on a range of other indicators – unemployment and other measures of slack in the labour market, latterly wage and domestic price inflation as well – that tell us something about how the supply side is behaving and can therefore make policy more robust to uncertainty about it.

But it's not possible to offset that uncertainty entirely. When the economy is subjected to a greater range of shocks it can be harder to identify their individual contributions. Even if the observed economy were measured with perfect accuracy our understanding of what's caused it – the precise behaviour of these underlying economic forces – would still be imperfect. And in the real world there's inevitably a degree of inaccuracy in economic measurement. Currently, if only for a short period of time, there's a little more uncertainty than usual about the behaviour of unemployment. Official estimates of wage growth have been volatile and other indicators have exhibited slightly lower (if still very elevated) rates of growth through much of this year.

As a general matter, one implication of this additional uncertainty is that, in response to any given shock, the reaction of policy is likely to be somewhat more delayed than in a world of perfect and complete information. It takes time to understand the forces driving the economy, particularly if one's having to rely on nominal variables like services inflation and wage growth, things that would normally be seen as late-cycle indicators. The same goes for uncertainty relating to the accuracy of the observable variables themselves.

One current example is wage growth. Given the volatility in the official estimates, and the disparity (such as it is) among the various indicators we have, it will probably require a more protracted and clearer decline in these series before the MPC can safely conclude that things are on a firmly downward trend.

Thank you.

I've received helpful comments from colleagues at the Bank of England. I'd like to thank Rich Harrison, Tom Key, Jack Page, Kate Reinold, and especially Fabrizio Cadamagnani and Lindsey Rice-Jones for their help in preparing the speech. The views expressed are my own and do not necessarily reflect those of the Bank of England.