Institute for Fiscal Studies

IFS Report R313

Max Warner
Ben Zaranko

The fiscal implications of public service productivity

Economic and Social Research Council

The fiscal implications of public service productivity

Max Warner
Ben Zaranko

Copy-edited by Rachel Lumpkin

Published by **The Institute for Fiscal Studies**

ISBN 978-1-80103-181-3

The authors gratefully acknowledge the support of the ESRC Centre for the Microeconomic Analysis of Public Policy (ES/T014334/1) and the UKRI project 'Productivity, Wages and the Labour Market' (ES/W010453/1).

1. Introduction

The challenging fiscal situation facing the next government means that both Labour and the Conservatives are looking for ways to improve public services without large increases in spending. Improving the productivity of public services is often proposed as a way to achieve this. This follows the big fall in measured productivity of government services – particularly in the NHS (Warner and Zaranko, 2023) – since the start of the pandemic. In the March 2024 Budget, the Chancellor, Jeremy Hunt, launched a new 'Public Sector Productivity Plan' and argued that 'the way to improve public services is not always more money or more people – we also need to run them more efficiently' (HM Treasury, 2024). Mr Hunt set out £4.2 billion of funding to improve the productivity of public services, including £3.4 billion for the NHS.

Improving the productivity of public services is desirable, as it means that the government can deliver more or better services with the same resources, or deliver the same services with fewer resources. Finding a way to do more with the same amount of resources is particularly important given the current combination of poor public service performance and a challenging fiscal climate (Emmerson et al., 2024). For the latter, weak growth and high debt interest payments are constraining the government's ability to increase public service funding without accompanying tax rises. Boosting public sector productivity would be one way to deliver the service improvements desired by governments without a need for extra spending.

Indeed, Mr Hunt was explicit about these motivations for higher productivity in his Budget speech: 'if we [deliver] cash-releasing savings, as we are committed to doing, it will be possible to live with more constrained spending growth without cutting services valued by the public' (HM Treasury, 2024).

In this report, we consider the current state of public service productivity and make two arguments about the potential fiscal implications of any future improvements. First, higher productivity – as measured by the Office for National Statistics (ONS) – would mean that public services have become more proficient at translating inputs (such as staff) into outputs (such as healthcare appointments). But funding might still need to rise to maintain the quantity and quality of services provided, if the prices of those inputs (such as wages) are growing more quickly than productivity.

Second, there are good reasons to suppose that an improvement in public sector productivity would not lead to 'cash-releasing' fiscal savings – in the short term, at least. In part, this is because higher productivity will result in higher public service performance. Being able to save

funding from higher productivity requires the government to 'bank' that money rather than allow service performance to improve, which would be challenging given the current poor performance of many public services.

In sum, productivity improvements might allow for spending to grow *at a slower rate than would otherwise be the case*, or might alleviate the pressure for future budget top-ups, but would not automatically unlock cash-releasing savings. In other words, public sector productivity improvements would be 'good' for the public finances, and undoubtedly something we should aim for. The optimistic interpretation of the post-pandemic productivity dip is that there is plenty of potential for catch-up growth. But there is no direct read across from official measures of public service productivity to the level of public spending – especially when there is such pressure for public service performance to improve – and governments would be wise not bank on their ability to achieve substantial cash-releasing savings.

2. How should we think about public service productivity? How is it measured?

Productivity captures the relationship between the outputs or outcomes of a system (e.g. the number of pupils taught and their grades) and the inputs (e.g. the number of teachers and schools). Higher outputs or outcomes for the same level of inputs, or the same level of outputs or outcomes for a lower level of inputs, constitute higher productivity.

Although simple in theory, measuring productivity can be challenging in practice. For the wider economy, the outputs of firms can be valued using market prices. These prices capture the value that consumers themselves place on different outputs. But most public services are provided free at the point of use, and so there are no prices with which to value their outputs. This also makes it more difficult to place a value on the outcomes produced by those services. Ideally, we might value a hospital admission by the amount by which it improved the patient's health, or a year of education by the benefits it has for a child's development and future life trajectory. In practice, this is rarely feasible.

One simple option, adopted by some statistical offices (Office for National Statistics, 2022a), is to define public service output as simply being equal to the input. The UK's ONS takes a more sophisticated approach and produces measures of productivity in public services based on the quantity of outputs provided, with some adjustments for the quality of those outputs.

This standard measure, produced by the ONS, is what the government has used when discussing the potential gains from improving productivity. This measure is based on the recommendations of the Atkinson Review (Atkinson, 2005), but at the Chancellor's request the ONS is currently working to improve the measure.

The ONS measures outputs in different ways for different public services. For some services – such as much of health and education – the ONS measures outputs in terms of the quantity of services provided (e.g. the number of medical treatments or number of hours of teaching delivered) adjusted for some measures of quality (e.g. the waiting time for treatment or the grades of students). This is as close to the ideal of valuing outputs by their outcomes as is currently possible.

But for other services – such as social security administration – it is difficult to measure the quality or outcomes of services, and so the ONS measures just the quantity of services provided without any adjustment for quality or outcomes. For some other services, it is conceptually difficult to define the outputs or outcomes of the service (e.g. the output of defence spending could be defined as the ability of the UK to conduct overseas military operations, the perceived or actual safety of UK citizens from armed conflict, or in numerous other ways). It can also be difficult to measure the outputs or outcomes of some services (e.g. a reduction in the amount of recorded crime could occur because the police have been effective in catching and deterring criminals, or because victims have stopped reporting crimes to an inefficient and unresponsive police force). In these cases, the ONS takes the simple approach discussed above and assumes that the outputs equal the inputs, which mechanically means there can be no changes in the measured productivity for these public services.

The ONS also measures the quantity of inputs used to produce each public service. These inputs are split into staffing (e.g. the number of teachers), the quantity of intermediate goods and services used (e.g. the number of textbooks) and use of the capital stock (e.g. the use of school buildings). Where possible, the quantity of inputs is measured directly but in many cases the quantity is measured indirectly using total expenditure on each input (e.g. spending on staffing) adjusted by a relevant deflator (e.g. the average wages of the staff or a similar group). Note that the ONS is here trying to measure something closer to the actual quantity of inputs used to produce public services, not the funding used to pay for them (something we return to below).

Dividing the output measure by the input measure gives the ONS public service productivity measure. Throughout this report, we refer to this ONS productivity measure as ‘input productivity’ because it is based on the quantity of inputs used to produce public services.

3. What has happened to public service productivity since the start of the pandemic?

The current focus on public service productivity comes following a large fall in measured productivity since the start of the COVID-19 pandemic. Figure 1 shows the ONS measure of input productivity between 1997 and 2023 (relative to 2019). Note that for 2022 and 2023, this measure does not include quality adjustments for public service outputs.

Figure 1. ONS measure of public service productivity (input productivity)

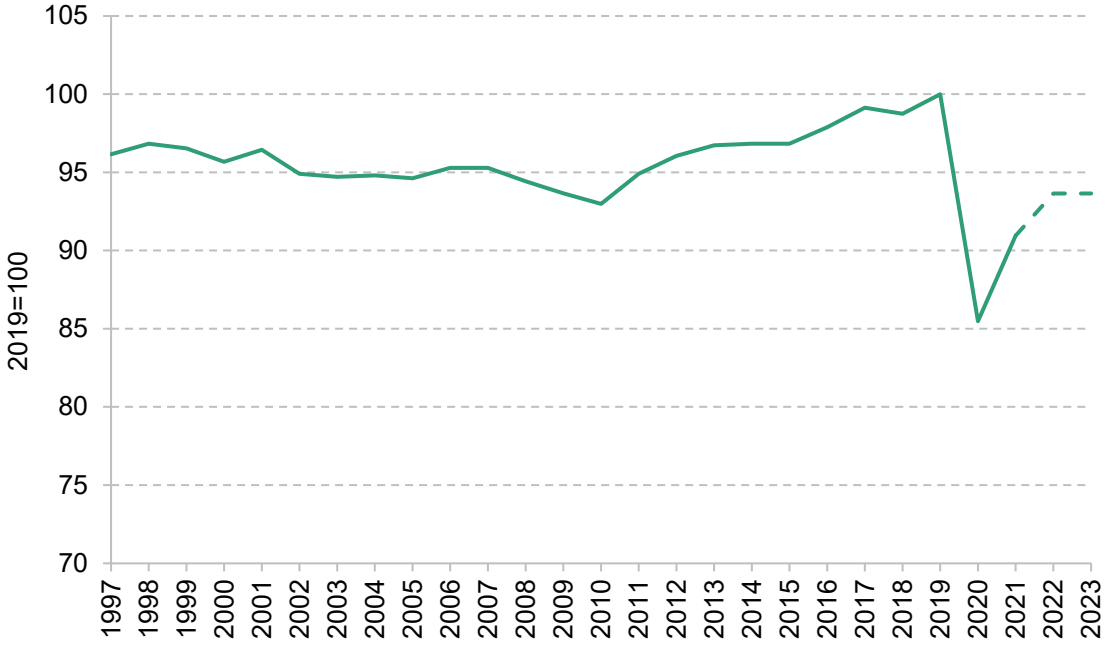

Note: Data for 2022 and 2023 (dashed lines) measure public service outputs without any quality adjustments

Source: ONS, Public service productivity, quarterly, UK, https://www.ons.gov.uk/economy/economicoutputandproductivity/publicservicesproductivity/datasets/publicserviceproductivityquarterlyuk.

Measured public service productivity declined slightly between 1997 and 2010, before rising between 2010 and 2019. We discuss these pre-pandemic trends in more detail in Section 4. Here we focus on what has happened since the start of the pandemic. In 2020, measured public sector productivity fell by 14.5%. Productivity then started to recover, rising by 6.4% in 2021 and 3.0% in 2022. However, there was no change in productivity between 2022 and 2023, suggesting that the recovery had stalled. Therefore, although productivity has partially recovered from its big decline in 2020, it remained 6.3% below 2019 levels in 2023. This is equivalent to losing the decade of public service productivity growth experienced prior to the pandemic.

There are a whole range of reasons for this reduction in public service productivity, and a detailed discussion is beyond the scope of this report. During the worst of the pandemic in 2020, there were many structural changes to the provision of services – such as reductions in school attendance, and cancellation of routine NHS activity to free up capacity for COVID-19 patients – that mechanically act reduce to productivity. But these direct impacts of the pandemic cannot explain the continued failure of productivity to return to pre-pandemic levels in 2022 and 2023. In the case of the NHS, we have argued that falls in productivity may be explained by changes in the mix of inputs, in particular the prioritisation of frontline staffing relative to spending on capital and management (Warner and Zaranko, 2022, 2023). NHS England have recently published their own analysis of NHS productivity, estimating that productivity in the acute hospital sector was 11% lower in 2023–24 than pre-pandemic (NHS England, 2024). This is measured differently to the ONS measure of productivity and is closest to the measure of funding productivity we discuss later. NHS England argues that lower productivity has been driven by a range of factors, including industrial action and temporary staffing costs.

It may be that the productivity shortfall relative to pre-pandemic levels provides grounds for cautious optimism, if we think it means there is considerable scope for catch-up growth. In other words, we don’t need to be at the cutting-edge of the productivity frontier; we just need to get back to where we were half a decade ago. For the remainder of this report, we abstract away from this debate, and abstract away from the huge oscillations in inputs and outputs that have occurred in recent years. We focus instead on the bigger-picture issues and trends at play, which we might expect to apply in a ‘normal’ world. From here on in, therefore, the analysis will largely focus on the two decades of pre-pandemic public sector productivity performance.

4. Higher measured public service productivity does not automatically mean that services become cheaper

As discussed earlier, the ONS measure of productivity is based on the quantity of inputs used to produce public services. There are good reasons to think about productivity this way, but changes in this measure of productivity cannot automatically be translated into changes in spending or funding required. This is because spending on public services depends not only on the quantity of inputs used, but also the price of inputs.

In this section, we therefore introduce another measure of productivity, which we call 'funding productivity'. This measure relates the outputs of public services to the funding used to provide them. Box 1 provides definitions of all terms used throughout this report.

A simple example

To understand why funding and input productivity may differ, and may have different implications, consider the following simple example. A GP practice has 10 GPs who together can deliver 100 appointments in a day. Input productivity is measured as output divided by input, so input productivity is 10 appointments delivered per GP (here we focus on labour input alone for simplicity). After some improvements in organisation, the same number of GPs can now deliver 105 appointments in a day. Input productivity has increased to 10.5 appointments per GP, and so is 5% higher than before. This is, in essence, what the ONS measure of public service productivity captures.

So far everything has been in quantities – the number of appointments divided by the number of GPs. But for the public finances, it is spending that matters. Let's say, again for simplicity, that the daily wage of a GP is £100. We define funding productivity as output divided by spending. Prior to the change, 100 appointments could be delivered by 10 GPs costing £100 each, so total funding was £1,000 per day. Funding productivity is therefore 0.1 appointment per £1 spent (100 appointments/£1,000 total funding). If wages remain constant, funding productivity after the organisational improvement will be 5% higher than before, the same as the increase in input productivity (105 appointments/£1,000 total funding = 0.105 appointments per £1 spent).

Box 1. Definitions of measures used

Output. The quantity of outputs produced by public services, as measured by the ONS. For some public services this output is quality-adjusted; for some others it is measured on an input=output basis.

Input. The quantity of labour, intermediate goods and services and capital inputs used by public services, as measured by the ONS.

Real funding. A measure of government day-to-day spending on public services. This is defined as being equal to total managed expenditure minus public sector net investment, social security spending and debt interest spending. This is deflated using the GDP deflator, the measure of economy-wide domestic inflation typically used for public finance calculations.

Input productivity. How much output public services deliver relative to the quantity of inputs. Defined using the above measures as output/input.

Implied input prices. The real prices of inputs. Defined using the above measures as real funding/inputs. This captures how the prices of inputs have changed relative to average prices in the economy.

Funding productivity. How much output public services deliver relative to their real day-to-day spending. Defined using the above measures as output/real funding

But imagine that the wages of GPs increased at the same time as their productivity. Let's assume that the daily wage of a GP increases by 10% to £110, perhaps in order to keep pace with rising wages in the other economic sectors that represent GPs' outside option. Then even with the improvement to input productivity from the organisational change, funding productivity is now lower than it was previously. That's because the total cost of providing the service (10 GPs * £110 = £1,100, a 10% increase) has increased by more than the increase in output (5% more appointments). As a result, funding productivity is now 0.095 appointments per £1 spent (105 appointments/£1,100 total funding), 5% lower than before, despite the 5% increase in input productivity.

Therefore, even if public services become more productive in the way measured by the ONS (input productivity), this does not automatically mean that they become cheaper to provide. Changes in the costs of inputs may offset any increases in how productively those inputs are being used. The key point is that, in general, if the prices of inputs rise by more than the increase in input productivity, funding productivity will decline.

Funding and input productivity between 1997 and 2019

To illustrate the importance of this distinction in practice, we can produce our own simple measure of funding productivity. We start with the same measure of public service output used to produce input productivity. But instead of dividing this by a measure of inputs to public services, as the ONS does to produce input productivity, we can instead divide it by (a broad measure of) real day-to-day spending on public services. This gives a measure of public service output per pound of resource funding received (i.e. funding productivity). Note that we are here excluding capital spending (i.e. spending on investment in machinery, equipment and buildings). This is because we are interested in how funding in a particular period translates into output *in the same period*. Capital spending is unquestionably important for determining the productivity of public services but is better thought of as investments in making services more productive in future, over an extended period. Contemporaneous, single-period comparisons of capital spending and public service outputs therefore don't make much sense. For more detail on definitions used throughout this piece, see Box 1.

Figure 2 shows how the different measures we have discussed changed on average between 1997 and 2019. For simplicity we focus on the pre-pandemic period, as the COVID-19 pandemic caused large changes in all measures (and complicated the collection and interpretation of those measures). It remains to be seen whether these changes will prove permanent, or whether some degree of catch-up growth will prove possible.

Starting first with the standard ONS measures, public services have increased their outputs by an average of 2.3% per year. Over the same period, their inputs have increased by 2.2% per year, which (given rounding) implies a 0.2% increase in input productivity per year. Over the two decades prior to the start of the pandemic, public service productivity therefore increased on average, albeit slowly (differences between subperiods are discussed below). In other words, to provide the same quantity and quality of public services required fewer inputs over time. For comparison, the ONS measure of total factor productivity (Office for National Statistics, 2022b) in the private sector (closest to, but slightly different from, the measure of input productivity of public services) grew by 0.8% per year over the same period (i.e. more quickly).

Figure 2. Components of UK public service productivity growth between 1997 and 2019

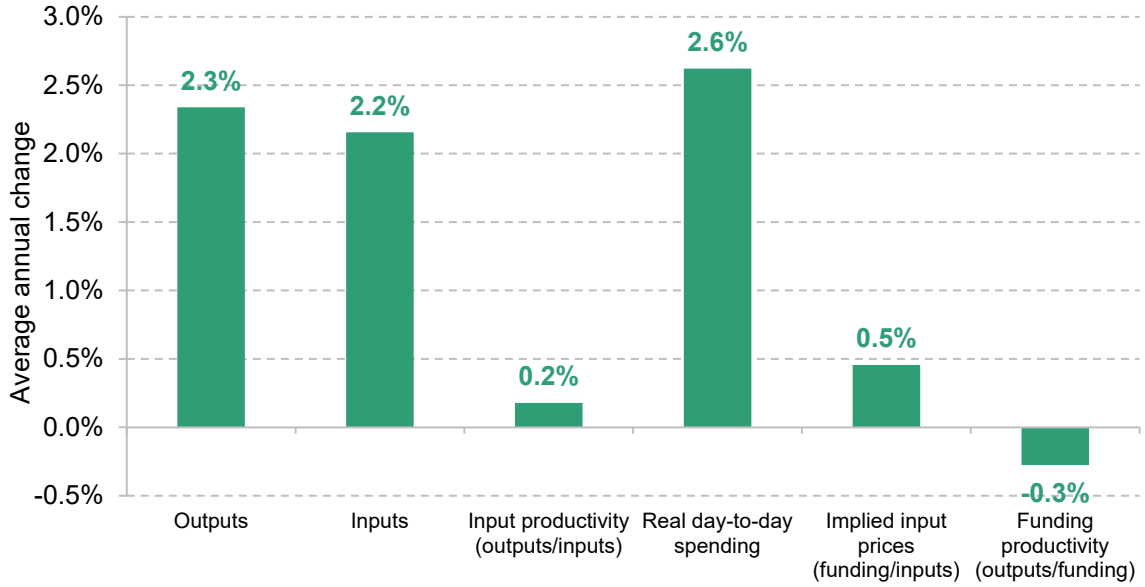

Note: See Box 1 for definitions of terms. Inputs and outputs are measured using calendar years, while real spending is measured using financial years. Figures may not sum due to rounding.

Source: Authors' calculations using: ONS, Public service productivity, quarterly, UK, https://www.ons.gov.uk/economy/economicoutputandproductivity/publicservicesproductivity/datasets/publicserviceproductivityquarterlyuk; IFS, Spending Composition Spreadsheet, https://ifs.org.uk/taxlab/taxlab-data-item/ifs-spending-composition-sheet; HM Treasury, GDP deflators at market prices, and money GDP March 2024, https://www.gov.uk/government/statistics/gdp-deflators-at-market-prices-and-money-gdp-march-2024-quarterly-national-accounts.

But while inputs to public services increased by 2.2% per year on average, our best measure of public service *funding* increased by 2.6% per year in real terms. This suggests that input prices – such as wages and equipment costs – rose by 0.5% per year in real terms (i.e. over and above the average increase in economy-wide prices). Consistent with this implied rise in input prices, ONS data (Office for National Statistics, 2023) on earnings show directly that public sector wages for full-time workers – one important input price – rose on average by 0.9% per year in real terms over this period.[1]

Because real funding grew faster than outputs (2.6% versus 2.3%), our measure of funding productivity fell by an average of 0.3% per year. Equivalently, input productivity grew more slowly than input prices (0.2% versus 0.5% per year). As a result, these figures suggest that public services became better at translating inputs into outputs more productively, but those services still became more costly to provide over time, because the cost of those inputs grew more quickly than did their productivity. This is evidence that the distinction between funding

[1] This uses the GDP deflator, which is the measure of economy-wide domestic inflation used for public finance calculations, rather than the Consumer Price Index, which is typically used for real earnings calculations.

and input productivity makes a first-order difference to the fiscal implications of higher productivity. It is important to note, however, that 'more expensive to provide' does not equate to 'less affordable'. Real GDP grew by an average 1.9% per year over this period. Public services became more costly to provide, but our ability to meet those costs (out of GDP) increased more quickly, making those public services more affordable. This is relevant to our discussion of the Baumol effect below.

Differences in funding and input productivity over time

Table 1 repeats our analysis separately for 1997 to 2009 and 2009 to 2019. These periods correspond to the New Labour governments and Conservative-led governments prior to the COVID-19 pandemic, respectively. Consistent with Figure 1, there was a substantially different pattern in overall productivity between these two periods.

Between 1997 and 2009, both the inputs and outputs of public services grew rapidly, but outputs grew more slowly than inputs, which meant that input productivity declined by 0.2% per year. Funding grew faster than inputs, implying a 0.6% average annual increase in real input prices, and so funding productivity fell by 0.8% per year. Therefore, although public services were delivering much greater output over this period, both input and funding productivity fell.

Table 1. Components of UK public service productivity growth between 1997 and 2019

Measure	Average growth rate between 1997 and 2019	Average growth rate between 1997 and 2009	Average growth rate between 2009 and 2019
Outputs	2.3%	3.4%	1.1%
Inputs	2.2%	3.6%	0.5%
Input productivity (outputs/inputs)	**0.2%**	**−0.2%**	**0.7%**
Real day-to-day spending	2.6%	4.2%	0.7%
Implied input prices (spending/inputs)	0.5%	0.6%	0.2%
Funding productivity (outputs/spending)	**−0.3%**	**−0.8%**	**0.4%**

Note: See Box 1 for definitions of terms. Inputs and outputs are measured using calendar years, while real spending is measured using financial years.

Source: Same as Figure 2.

Between 2009 and 2019, both the inputs and outputs of public services grew much more slowly than over the previous decade. But input productivity rose by 0.7% per year because outputs grew at more than double the rate of inputs. Input prices only grew slowly in real terms (0.2% per year) – reflecting public sector pay restraint during this period (Cribb and O'Brien, 2024) – which meant that funding productivity grew by 0.4% per year.

Differences in funding and input productivity across public services

Table 2 repeats our analysis for two large public services: health and education. In both cases, implied input prices have risen substantially over time in real terms, meaning that improvements in input productivity translate into much smaller gains in funding productivity. For education, like public services as a whole, measured input productivity rose over this period (by 0.1% per year), but input prices grew more quickly (0.5% per year), suggesting that funding productivity fell (by 0.3% per year). For healthcare, measured funding productivity grew (by 0.3% per year) because the growth in input productivity (0.9% per year) was greater than the growth in input prices (0.6% per year).

Table 2. Components of UK public service productivity growth between 1997 and 2019 in healthcare and education

Measure	Healthcare	Education
Outputs	4.4%	1.6%
Inputs	3.5%	1.5%
Input productivity (outputs/inputs)	**0.9%**	**0.1%**
Real day-to-day spending	4.1%	1.9%
Implied input prices (spending/inputs)	0.6%	0.5%
Funding productivity (outputs/spending)	**0.3%**	**−0.3%**

Note: Figures denote the average growth rate between 1997 and 2019. See Box 1 for definitions of terms. Inputs and outputs are measured using calendar years, while real spending is measured using financial years.

Source: Same as Figure 2 and ONS, Public service productivity estimates: healthcare: UK, https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthcaresystem/datasets/publicserviceproductivityestimateshealthcareuk; ONS, Public service productivity estimates: education, https://www.ons.gov.uk/economy/economicoutputandproductivity/publicservicesproductivity/datasets/publicserviceproductivityestimateseducation.

That a similar pattern occurs for these two large public services, where outputs and inputs are much better measured than for other public services, also suggests that our primary analysis is not driven by the worse measures of inputs and outputs used for other public services.

The Baumol effect

One way to think about these trends is through the lens of what is known as the Baumol effect, first described by William Baumol and William Bowen in the 1960s (Baumol and Bowen, 1965). The core of the argument is that wages (and other input prices) in all sectors tend to rise with economy-wide productivity growth, but that productivity growth is unevenly distributed across parts of the economy. In other words, wages rise even in the sectors that have experienced minimal productivity growth, because firms in these low-productivity growth sectors have to compete for workers with the sectors where productivity has risen. The result is that input costs in those low-productivity sectors tend to rise faster in real terms and relative to the rest of the economy.

Baumol and Bowen originally discussed this effect in the context of the performing arts. The classic example is that of a performance of Beethoven's String Quartet No. 14 (Helland and Tabarrok, 2019). In 1826, when it was first performed, it took four people 40 minutes to perform. In 2024, it still takes four people 40 minutes to perform. The input productivity of the string quartet is unchanged. But they are paid much more in 2024, owing to productivity improvements (and therefore wage increases) in other sectors, with which wages in the classical music sector must compete. Therefore, the funding productivity is much lower.

The public sector includes many labour-intensive services where we might similarly expect the Baumol effect to apply. Education is a good example of a setting where substantial productivity growth – at least in terms of raw outputs – is unlikely: it takes a teacher just as much time today to mark an essay as it did in the past, and school class sizes are much the same as in decades past. The data reported in Table 2 are consistent with this: input productivity in education increased by just 0.1% per year between 1997 and 2019. Baumol's theory would suggest that input prices in sectors such as education would rise even with such minimal improvements in productivity. This is indeed what we see: implied input prices in education grew by 0.5% per year between 1997 and 2019. The resulting reduction in what we term 'funding productivity' is consistent with the Baumol effect and represents an increase in the relative prices of services such as education over time.

If we look at the public sector as a whole, we similarly see trends consistent with the Baumol theory. Between 1997 and 2019, input productivity in the public sector grew more slowly than did (the roughly comparable) total factor productivity in the private sector (0.2% per year versus 0.8% per year). But real average wages of full-time employees (a proxy for the input price of labour) grew at the same average rate (0.9% per year) in the public and private sectors over that period. Public services therefore became more expensive to provide over time. But as discussed above, they simultaneously became more affordable over the two decades prior to the pandemic, because our ability to pay for them – as proxied by real GDP – grew by even more. The problem in recent years has been the combination of stalling (or falling) productivity alongside stagnant economic growth.

5. Higher public service productivity might not produce fiscal savings in the short term

As discussed above, aiming to improve public sector productivity is entirely sensible. Relative to a counterfactual with no productivity improvements, we would expect higher productivity to reduce the amount that needs to be spent, to improve the output and quality of public services, or both. Even if spending still needs to rise over time (due to the factors discussed above), public service productivity improvements might allow it to increase more slowly and thus provide savings relative to the counterfactual. A simplified version of the relationship between productivity, spending and (quality-adjusted) output is illustrated in Figure 3.[2]

On the horizontal axis is spending on public services, while on the vertical axis is the quality-adjusted output of public services. The green line represents the starting relationship between the two: higher spending leads to a higher quality and quantity of public service output. An improvement in public service (funding) productivity shifts this line anticlockwise, to the yellow line. This means that relative to the starting point A, it is now possible to spend the same and deliver a better-quality output (point B), or it is cheaper to provide the same quality-adjusted output (point C).

Points B and C represent two potential responses to the increase in productivity. In the former case, all of the increase in productivity is used to increase the quality-adjusted output of public services. This might manifest in greater output (e.g. more patients treated by the NHS) or greater quality (e.g. shorter waiting times or better educational attainment). In the latter case, all of the increase in productivity is used to reduce funding while maintaining the quality-adjusted output of public services.

[2] Here we abstract away from the public services where there is no quality adjustment in the ONS measure.

Figure 3. Stylised relationship between spending and quality-adjusted output

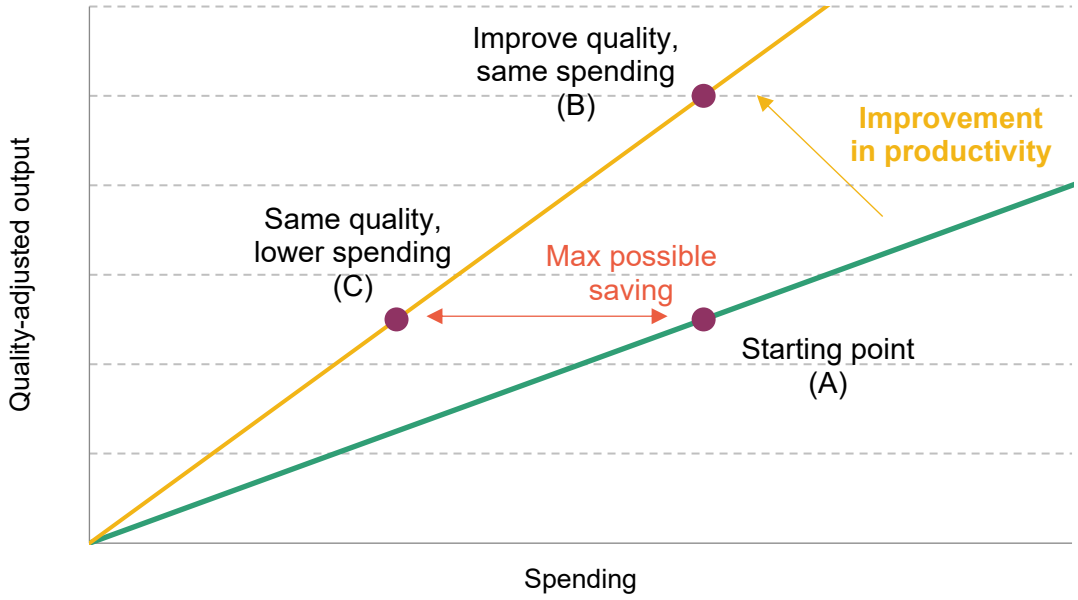

Note: This is a stylised relationship. In practice, the relationship between spending and quality-adjusted output is likely non-linear and differs between different public services.

The key point is that the government has a range of options – it has a choice over how to respond to an increase in productivity. It could spend the same and increase quality-adjusted output. It could reduce funding and hold quality-adjusted output constant. Or the government might decide that since the returns to spending are now greater, it wishes to spend more (and deliver an even bigger increase in quality-adjusted output). This is fundamentally a political choice. But in present circumstances, there are good reasons to think that higher public service productivity is more likely to show up as improved quality-adjusted output rather than big reductions in spending.

First, public service performance and quality are currently low relative to both historic performance and the government's own targets. This is the case for the NHS (Warner and Zaranko, 2024), prisons (Hoddinott et al., 2023a), courts (Hoddinott et al., 2023b) and many other public services. The low baseline of current performance and the numerous targets to improve performance mean that it is almost certain political pressure will mean much of any productivity improvements go towards improving the quality of public services rather than freeing-up funding.

Second, unlocking cash-releasing savings might prove easier said than done. Trivially, it's true that many ways to improve productivity, such as investing in technology, require upfront investment to deliver future benefits. But beyond that, the ways in which public sector organisations are incentivised to improve productivity may reduce the amount of funding that

can be saved from such improvements. The Health Secretary, Victoria Atkins recently announced, for example, that there would be new incentives to reward NHS trusts that improve productivity (Illman, 2024). She stated that 'I want to see providers retain the surpluses generated through productivity'. As a way of providing NHS leaders with incentives to pursue measures that might enhance productivity, this makes perfect sense. But by definition, this will limit any fiscal savings: if organisations retain the surplus from improvements, spending does not fall.

In practice, then, we might expect most of the productivity improvements to be absorbed by quality and output improvements in the short term.

The longer-term effects might be different. Relative to the counterfactual of *no productivity growth*, we would expect productivity improvements to ease spending pressures. Spending might still need to rise (for the reasons discussed earlier), but productivity improvements allow it to rise more slowly. Consider, for example, the NHS England Long Term Workforce Plan, which would see the NHS workforce grow from around 1.5 million now to between 2.3 and 2.4 million by 2036–37 in order to meet rising demands upon the service. The plan is predicated on labour productivity improvements of 1.5% to 2% per year. Without such improvements, an even bigger increase in staff numbers would be required, and the plan would become even more expensive to deliver. Or, more optimistically, if productivity improvements outpace expectations, then staffing numbers might not need to grow by quite as much, lowering the cost. This comes back, however, to the political trade-off between spending and quality-adjusted output discussed above. To reduce spending would still require a deliberate choice not to allow quality-adjusted output to improve further.

6. Conclusion

Improvements in public service productivity are unquestionably a good thing. Finding ways to deliver more and better public services with the same resources is obviously something that governments should pursue. Productivity improvements would give the government a choice over whether to increase the quality and output of public services, or to reduce funding while maintaining quality and output – a good choice to have.

But there seem to be misconceptions about the fiscal implications of improvements to public service productivity. In this piece, we seek to address two. First, improvements in public service productivity, as measured by the ONS, do not necessarily mean that the amount of funding required to deliver a given level of output will fall over time, if – as economic theory would predict – the price of public service inputs (such as staff wages) rise over time. Productivity improvements might allow for spending to grow *at a slower rate than would otherwise be the case* and are still therefore 'good' for the public finances, but there is no direct read across from official measures of public service productivity and the level of public spending. Second, and relatedly, there are good reasons to expect improvements in public service productivity to manifest as quality improvements rather than cash savings – at least in the short term. As a result, relying on cash-releasing savings from public service productivity improvements alone to get us out of our current fiscal predicament would be unwise.

References

Atkinson, A. (2005), *Atkinson Review: Final report. Measurement of Government Output and Productivity for the National Accounts*. New York: Palgrave Macmillan. https://webarchive.nationalarchives.gov.uk/ukgwa/20160105160709/http://www.ons.gov.uk/ons/guide-method/method-quality/specific/public-sector-methodology/articles/atkinson-review-final-report.pdf

Baumol, W. J. and Bowen, W. G. (1965). On the performing arts: the anatomy of their economic problems. *American Economic Review*, 55 (1/2), 495–502. https://www.jstor.org/stable/1816292

Cribb, J. and O'Brien, L. (2024). Recent trends in public sector pay. IFS Report R305. https://ifs.org.uk/publications/recent-trends-public-sector-pay

Emmerson, C., Farquharson, C., Johnson, P. and Zaranko, B. (2024). Constraints and trade-offs for the next government. IFS Report R295. https://ifs.org.uk/publications/constraints-and-trade-offs-next-government

Helland, E. and Tabarrok, A. (2019), *Why Are the Prices So Damn High? Health, Education, and the Baumol Effect*. Annals of Computational Economics, George Mason University, Mercatus Center, May. https://www.mercatus.org/students/research/books/why-are-prices-so-damn-high

HM Treasury (2024). Spring Budget 2024 Speech. https://www.gov.uk/government/speeches/spring-budget-2024-speech

Hoddinott, S., Davies, N., Fright, M., Nye, P. and Richards, G. (2023a), Performance Tracker 2023. Institute for Government. https://www.instituteforgovernment.org.uk/publication/performance-tracker-2023/prisons.

Hoddinott, S., Davies, N., Fright, M., Nye, P. and Richards, G. (2023b), Performance Tracker 2023. Institute for Government. https://www.instituteforgovernment.org.uk/publication/performance-tracker-2023/criminal-courts.

Illman, J. (2024). Atkins hints at 'new incentives' for trusts to hit productivity targets. HSJ. https://www.hsj.co.uk/finance-and-efficiency/atkins-hints-at-new-incentives-for-trusts-to-hit-productivity-targets/7036735.article

NHS England (2024). NHS productivity. https://www.england.nhs.uk/long-read/nhs-productivity/

Office for National Statistics (2022a). International comparisons of the measurement of non-market output during the coronavirus (COVID-19) pandemic.

https://www.ons.gov.uk/economy/economicoutputandproductivity/output/methodologies/internationalcomparisonsofthemeasurementofnonmarketoutputduringthecoronaviruscovid19pandemic

Office for National Statistics (2022b). Multi-factor productivity estimates. https://www.ons.gov.uk/economy/economicoutputandproductivity/productivitymeasures/datasets/multifactorproductivityexperimentalestimatesreferencetables

Office for National Statistics (2023). Annual Survey of Hours and Earnings time series of selected estimates. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/ashe1997to2015selectedestimates

Warner, M. and Zaranko, B. (2022). NHS funding, resources and treatment volumes. IFS Report R236. https://ifs.org.uk/publications/nhs-funding-resources-and-treatment-volumes

Warner, M. and Zaranko, B. (2023). Is there really an NHS productivity crisis? IFS Comment. https://ifs.org.uk/articles/there-really-nhs-productivity-crisis

Warner, M. and Zaranko, B. (2024). The past and future of NHS waiting lists in England. IFS Report R302. https://ifs.org.uk/publications/past-and-future-nhs-waiting-lists-england