National Foundation for Educational Research

Technical Report

NFER Teacher Supply Forecast and Simulation Model

National Foundation for Educational Research (NFER)

NFER Teacher Supply Forecast and Simulation Model: Technical Report

Jack Worth, Sarah Tang and Dawson McLean

Published in January 2026

By the National Foundation for Educational Research,

The Mere, Upton Park, Slough, Berkshire SL1 2DQ

https://www.nfer.ac.uk/

ISBN: 978-1-916567-47-4

How to cite this publication:

Worth, J., Tang S. and McLean, D. (2026). *NFER Teacher Supply Forecast and Simulation Model: Technical Report*. Slough: NFER.

Contents

Acknowledgements

We are very grateful for the grant funding provided for building this model by the Gatsby Charitable Foundation and the support throughout the project from Jenni French. Thank you very much to Sarah Tang for running the School Workforce Census (SWC) analysis and to Dawson McLean for designing, coding and testing the simulation model. We are also very grateful to the Teacher Workforce Model (TWM) team at the Department for Education (DfE) for the support they have given us in understanding the model.

NFER
National Foundation for Educational Research

1. Introduction

In 2025, the National Foundation for Educational Research (NFER) was funded by the Gatsby Charitable Foundation to develop a teacher supply forecast and simulation model. The aim of the model is to create a representation of the state-funded schoolteacher workforce in England and simulate its evolution over time according to recruitment, retention and career decision dynamics.

The model allows the recruitment and retention dynamics within the workforce to be influenced by economic trends and changes in key policy levers, such as annual pay increases, training bursaries and the design of retention payments. How teachers respond to changes in the economy and policy changes comes from parameters estimated in the best available research evidence.

The model captures the uncertainty that is inherent in making a forecast about the future. The decisions that synthetic teachers are assumed to make are governed by probabilities and the forecasts of the relevant factors (such as pupil number projections, economic trends) are assumed to vary and not be deterministic. The parameters that determine how responsive synthetic teachers' behaviour is to key relevant factors is also assumed to vary.

To allow the model to capture this uncertainty, while still deriving useful outputs that reflect the best current estimate, we run the model as a 'Monte Carlo' simulation. This means that we allow factors in the model to vary according to randomness but run the model lots of times to even out that randomness. This allows us to derive a central estimate that tells us the most likely forecast given our assumptions, while also providing an understanding of how uncertain that forecast might be and how much it might be expected to typically deviate from that central forecast. Section 7 below explains more about how the model handles this randomness.

The sections of this technical report about the model are arranged as follows. Section 2 describes how we estimate the baseline synthetic population of teachers. Section 3 explains how we model the career dynamics within the teacher workforce, incorporating teachers leaving and new teachers entering to derive the evolution of the synthetic population over time. Section 4 explains how the initial teacher training (ITT) targets are estimated while section 5 explains how policy costs are estimated. Section 6 explains what outputs the model produces and how it does so. Section 7 describes in more detail how the model captures uncertainty through the 'Monte Carlo' aspect of the model set up.

2. Synthetic population of teachers

The forecast and simulation model is **agent-based**, which means it models the career decisions of individual teachers and aggregates them to understand the trends and dynamics at the population level. The model is not based on real teachers, but a **synthetic population** of hypothetical teachers, represented in a dataset according to some of their key characteristics (such as what subject they teach, age, sex, years of experience, working pattern).

We used data from the 2023 School Workforce Census (SWC) – a Department for Education (DfE) dataset containing the characteristics of real teachers in state-funded schools in England in November 2023 – to calibrate the characteristics of the synthetic population. This means that the distribution of teacher characteristics in the synthetic population matches those in the real population, while not directly representing any real-world individual teachers.

The synthetic population is built up sequentially. The full-time equivalent (FTE) number of teachers by subject is the basis of the model, which is derived directly from the numbers published by DfE in its Teacher Workforce Model (TWM). This does mean the total number of teachers in the population differs from the total number of teachers employed in state-sector schools according to the SWC. For example, the TWM reports that there were 424,537 FTE teachers in 2023/24, whereas the SWC reports that there were 468,693. This difference arises for several reasons, including differences in institutional coverage (TWM only includes teachers in primary and secondary schools and not e.g. special schools, alternative provision) and teachers included (e.g. leaders who do not teach a subject are not included in the TWM numbers). Unless otherwise stated, the outputs from the model are always based on a model that is calibrated to the latest available TWM.

We used regression analysis to ensure that the characteristics of teachers in our synthetic population matched those in the SWC. First, we ran a logistic regression using SWC data to estimate the probability of a teacher having a full-time working pattern, given the subject they teach. We then applied this probability to each teacher in the base synthetic population. We then assigned them either full-time or part-time status according to a uniform random variable that is different for each synthetic teacher.

Next, we ran a logistic regression using SWC data to estimate the probability of a teacher being female, given the subject they teach and their working pattern, and ran the same process as for working pattern to assign each synthetic teacher to be either male or female. This process proceeded sequentially, assigning synthetic teachers characteristics including their age, years of experience, pay region (Inner London, Outer London, London Fringe or Rest of England), and subject specialism, based on all characteristics assigned up to that point. Each full time teacher is assigned an FTE of one, while part-time teachers are assigned the average FTE among part-time teachers of each subject, estimated from the SWC.

Using regression analysis for this process ensured that not only do the synthetic population characteristics match the population on average, but also the correlations between characteristics (e.g. female teachers are more likely to work part-time than male teachers) are retained in the synthetic population as they are in the real population.

Finally, we estimated each synthetic teachers' remuneration. First, we used a linear regression and data from the SWC to estimate teachers' gross FTE total pay, according to the characteristics listed above. Gross total pay includes teachers' salary and any additional payments by the school, such as teaching and learning responsibility (TLR) payments.

Second, drawing on a range of policy documents published by DfE we assigned each teacher any early career retention payments that they were eligible for as part of their overall remuneration and that are paid directly by central Government. These included the targeted retention incentive, levelling up premium, phased maths bursary, early career payment and teacher student loan reimbursement schemes (Department for Education, 2017; DfE, 2019, 2022, 2024). Each synthetic teacher's eligibility and payment amount was determined based on their subject, years of experience and academic year. For some early career payments, eligibility was also determined according to school type or area criteria. To model this we estimated the proportion of eligible schools within each pay region and assigned eligibility and payment amounts to individual teachers according these probabilities (using a set of random variables that are different for each synthetic teacher).

3. Modelling the dynamics

The process described above generates a synthetic teacher population for the baseline year in the simulation model. To generate a forecast of how the teacher population may evolve over time, we also model the dynamics within the workforce. Dynamics are modelled in several different ways, including those that are influenced by policy changes set by the model user.

These are described in turn below and are:

- Trainees entering the state-funded sector
- Returners and other new teachers entering the state-funded sector
- Teachers leaving the state-funded sector
- Characteristic changes for teachers who stay in the state-funded sector

3.1. Trainees entering the state-funded sector

Each year newly qualified teachers enter the state-funded sector directly from completing their teacher training. These trainees enter the synthetic population according to first the number of trainees, then are assigned characteristics in a similar way to how the baseline synthetic population is generated.

The number of entrants by subject from high potential (Teach First), undergraduate and assessment only ITT are taken directly from the TWM, with the numbers beyond the final year reported in the TWM assumed to be the same in subsequent years.

We calculate the number of new entrants from postgraduate teacher training by subject based on the figures reported by the DfE in the ITT Census. Unless otherwise stated, outputs from the model are always based on the latest available ITT Census data. For cohorts that are relevant to the model but where ITT Census has not yet been published, we forecast the number of trainees in each subject. The starting point for the forecast is the number of trainees in the most recent available year in the ITT Census (baseline year). The forecast for the number of postgraduate trainees estimated for each subject then deviates from that starting point according to the following factors:

- The difference between the subject-specific bursary for that year and the bursary in the baseline year, multiplied by the recruitment elasticity of the bursary. The elasticity is assumed to be a 2.9 per cent change in the number of recruits per £1,000 change in the bursary (National Audit Office, 2016; Worth and Hollis, 2021; Worth, Tang and Galvis, 2022; CFE Research and FFT Education Datalab, 2023; McLean, Tang and Worth, 2023). The bursaries for future years are set as an input by the model user.
- The difference between the forecast of the unemployment rate for that year and the unemployment rate in the baseline year, multiplied by the recruitment elasticity of the unemployment rate. The elasticity is assumed to be a random parameter with a normal distribution, an average of a 6 per cent change in the number of recruits per percentage point change in unemployment rate and a standard deviation of 1.3 (Worth, Tang and Galvis, 2022). Unless otherwise stated, unemployment rate data comes from the latest Office for Budget Responsibility (OBR) forecast. The unemployment rate is a random parameter with a normal distribution, a mean of the forecast made by the OBR and a standard deviation derived from analysis of the historical differences between OBR forecasts and outturn data (see section 7 for details).
- The growth in teacher pay between that year and the baseline year, over and above the growth in forecasted average earnings growth between that year and the baseline year. Because not all information on expected pay is available when trainees decide whether to enter training, we model the influence of the pay growth difference as comprised of half of the concurrent difference (i.e. relevant to that year) and half of the lagged difference (relevant to the previous year). This pay difference is multiplied by the pay elasticity of recruitment. The elasticity is assumed to be a random parameter with a normal distribution, an average of a 2 per cent change in the number of recruits per percentage point difference in pay and a standard deviation of 0.4 (Worth, Tang and Galvis, 2022). Teacher pay growth for future years are set as an input by the model user. Unless otherwise stated, average earnings data comes from the latest OBR forecast. Average earnings is a random parameter with a normal distribution, a mean of the forecast made by OBR and a standard deviation derived from analysis of the historical differences between OBR forecasts and outturn data (see section 7 for details).
- Random variation in the number of trainees due to forecast uncertainty. We compared outturn data on trainee numbers by subject with a forecast like that described in this section. We estimated the level of precision by taking the standard deviation of the differences between the historic forecasts and the outturns. We then apply the random noise to the trainee number forecast for each subject using a normal distribution with a

mean of zero and standard deviation of the estimated prediction uncertainty (see section 7 for details).

The forecasted number of trainees estimated according to the process described above is then adjusted to account for the proportion of trainees that do not achieve qualified teacher status and the proportion who are not expected to enter the state-sector. These parameters are subject specific and taken directly from the TWM.

We then use the adjusted number of trainees (from postgraduate training and other sources) to generate a synthetic population of trainee teachers that can be added to the synthetic population in the following year. Trainee teacher characteristics are applied sequentially using a similar process to that described above to generate the overall synthetic population. The regression coefficients required for doing this sequential process were estimated from a sample of trainees in the SWC. All trainees entering the synthetic population are assumed to have one year of experience and to be subject specialists in their ITT subject. We assign trainees a sex, age, working pattern (and FTE), region (one of the nine Government office regions) and pay region.

The regression approach estimates an initial probability of a trainee being in a particular region, which is then adjusted depending on the difference between the subject-specific bursary for that year and the bursary in the baseline year. We use estimates of how the regional distribution depends on changes in the bursary level to make these adjustments (see McLean, Tang and Worth, 2023 for details).

3.2. Returners and other new teachers entering the state-funded sector

Each year teachers who are not recent trainees also enter the state-funded sector. These include returners, teachers who transferred their QTS from another country and teachers who have not previously worked in the state-funded sector but qualified to teach at least a year before. The number of such teachers entering the synthetic population each year are taken directly from the TWM, with the numbers beyond the final year reported assumed to be the same in subsequent years.

We take a similar approach to adding this group of new teachers into the synthetic population as for teachers in the original synthetic population, by sequentially estimating their characteristics from regression models using SWC data. Similar to the process we use to determine characteristics for teachers in the synthetic population, we assign new non-trainee teachers a working pattern (and FTE), sex, age, years of experience, pay region, subject specialism and gross pay.

3.3. Teachers leaving the state-funded sector

A proportion of teachers in the synthetic population in a particular year leave the state-funded sector before the next year. To model this process, we assign each teacher in the synthetic population a leaver status (i.e. whether they leave or not). If they are a leaver then we remove them from the synthetic population in the next year, whereas if they stay then they form the basis of the subsequent year's synthetic population. Even though in real life returners are individuals who have previously left, the model treats leavers and returners as separate groups of synthetic teacher.

To estimate synthetic teachers' probability of leaving, we ran a logistic regression to estimate the likelihood of a teacher leaving depending on their characteristics, using data from the SWC. Each synthetic teacher is assigned a leaving probability according to their characteristics.

Before being used to assign a leaver status, this probability is adjusted according to two factors:

- First, we take an FTE-weighted average of the leaving probabilities across all teachers in each subject in the synthetic population and compare it with the subject-specific leaving rates assumed in the TWM. We then re-calibrate the probabilities based on an adjustment factor, so that the average of the leaving probabilities in the synthetic population match those in the TWM. The two differ slightly due to differences in institution coverage, different data approaches to preparing the SWC data and other factors. This step ensures that the base leaving probabilities in the synthetic population match those assumed in the TWM.
- The growth in total teacher remuneration between that year and the baseline year, over and above the growth in forecasted average earnings growth between that year and the baseline year. This is multiplied by the pay elasticity of retention. The elasticity is assumed to be a 1.5 per cent reduction (or increase) in the leaving rate per percentage point that the growth in total teacher remuneration is higher (or lower) than the growth in average earnings (DfE, 2020). Total teacher remuneration includes gross pay and any early career retention payments the teacher is eligible for, both of which can be determined by the model user. Unless otherwise stated, average earnings data comes from the latest OBR forecast. Average earnings is a random parameter with a normal distribution, a mean of the forecast made by OBR and a standard deviation derived from analysis of the historical differences between OBR forecasts and outturn data (see section 7 for details).

Using this adjusted leaving probability, teachers in the synthetic population are assigned a leaver status according to a uniform random variable that is different for each synthetic teacher.

3.4. Characteristic changes for teachers who stay in the state-funded sector

Some teacher characteristics can change from year to year, which we model every time we generate a new year of the synthetic teacher population. The age and years of experience of teachers that stay in the synthetic population increase by one each year and teachers retain the same subject and sex.

Working pattern and pay region can change from year to year, due to teachers moving from full time to part time (or vice versa) and teachers moving school. We model this by using SWC data to estimate transition probabilities. For every synthetic teacher that stays in the state-funded sector, we use a logistic regression model to estimate the probability that they have full time status and the probability that they are in a particular pay region the following year, based on their existing characteristics. Each synthetic teacher's working pattern and pay region status is then updated according to a set of uniform random variables that are different for each synthetic teacher.

From year to year, each synthetic teachers' pay is re-estimated according to their new characteristics as well as uprated according to the assumed annual growth in teacher pay up to that academic year, which is an input set by the model user.

4. Estimating training targets

The model then uses various pieces of summary data within the model to forecast ITT targets for future years. This uses exactly the same methodology as the DfE's TWM, drawing on the number of trainees (forecasted as described above) and summaries of the synthetic population in each year. It also includes estimates of teacher demand, which rely on estimates of changes in pupil numbers from DfE's pupil projections.

In the TWM, estimates of projected teacher demand depend on expected growth in pupil numbers. Projections of the number of primary and secondary pupils in the model are random variables with a normal distribution. Unless otherwise stated, the average is the latest available pupil projections data published by DfE. The pupil projection for primary is the sum of pupils age Under five to ten, regardless of reported school phase. The pupil projection for secondary is the sum of pupils age 11 to 15, regardless of reported school phase, plus an assumption for the number of 16- and 17-year-olds[1]. The standard deviation of the random variable comes from our analysis of historic deviations between pupil projection forecasts and actual outturns.

The relationship between the change in pupil numbers and the change in teacher demand uses the same assumptions as used by DfE in the TWM.

5. Estimating costs

One of the model outputs is estimations of policy costs. These include costs paid by schools: gross teacher salary costs, employer national insurance contributions (NICs) and employer pension contributions. The estimates also include policy costs borne by central Government: training bursaries, early career retention payments and estimated central Government portions of teacher training costs.

Because the size of the teacher population in the TWM differs from that in the SWC, we rely on SWC estimates as the base for our estimates of total salary costs. For the baseline year we multiply the number of FTE teachers in primary and secondary schools from the published SWC

[1] The assumptions are that the number of 16 years olds in state secondary schools is 35 per cent of the previous year's 15-year-olds and the number of 17 years olds in state secondary schools is 35 per cent of the number of 15-year-olds from two years prior to that. The assumption of 35 per cent comes from DfE's participation in education, training and employment age 16 to 18 statistics.

data by the average FTE salary reported in the published SWC data. For subsequent years, we then uprate the average salary according to the pay growth assumed in the model and the number of teachers according to the change in the number of synthetic teachers. NICs and pension costs are estimated similarly.

Bursary costs are estimated by multiplying bursary levels by the number of trainees. This is likely to be a reasonable estimate of policy cost but may differ from actual spend due to some trainees being eligible for a (higher paying) scholarship and other trainees being ineligible for a bursary. Similarly, early career retention payment eligibility is estimated based on assumptions available in policy documents, but the overall policy cost may differ from our estimates due to exact eligibility criteria and claim rates. An estimate of the central Government portion of teacher training costs is taken from Allen *et al.*, (2016) and uprated for inflation, but may not reflect current policy costs due to changes in that portion over time. For these reasons, cost estimates should be treated as indicative.

6. Model outputs

The model outputs consist of summary data taken from the model. All outputs include the average among all of the model 'Monte Carlo' runs, the standard deviation and ventiles (i.e. the outcomes from all the model runs ranked and then split into 20 equal groups). These outputs include the number of trainees by subject and year, the respective targets and the proportion of the recruitment target met. Outputs also include summary data from the synthetic population by subject and year, including number of teachers, breakdowns by age, sex and experience level and the number of leavers. As explained above, the model also produces outputs on associated policy costs. The outputs can be compared across different scenarios run by the model user.

7. Modelling uncertainty

In sections above we explain how we use random variables rather than deterministic data on some forecasts where uncertainty is to be expected. For example, while the OBR produces a single forecast for the unemployment rate, based on past forecast reliability is it is reasonable to expect that it may deviate from that.

Therefore, we treat the forecasts of pupil number growth, unemployment rate and average earnings growth as random variables in the model. These are modelled as having a normal distribution with a mean of the published forecast and a standard deviation derived from our own analysis of the difference between historic forecasts and subsequent outturns. The standard deviations for each of these variables is summarised in Table 1 below.

Table 1 Forecast uncertainty related to model factors

Factor	Years after the baseline			
	Year 1	Year 2	Year 3	Year 4
Pupil numbers (primary)	0.6%	1.0%	1.3%	1.2%
Pupil numbers (secondary)	0.4%	0.6%	0.9%	1.2%
Unemployment rate	0.1pp	0.5pp	0.9pp	0.9pp
Average earnings growth	0.8pp	1.6pp	2.1pp	2.4pp

Source: NFER analysis of DfE and OBR data.

As explained above, we also assume random variation in the number of trainees due to forecast uncertainty. We compared outturn data on trainee numbers by subject with a forecast like that described in section 3, combining trainee numbers from prior years, information on what policy changes there have been and associated elasticities. We estimated the level of precision by taking the standard deviation of the differences between the historic forecasts and the outturns. We then apply the random noise to the trainee number forecast for each subject with a mean of zero and standard deviation of the estimated prediction uncertainty. The levels of forecast uncertainty are 21% in the first forecast year; 24% in year 2; 28% in year 3; 31% in year 4.

The model is run as a 'Monte Carlo' simulation, with the model run lots of times. The aim of running the model several times is to gain a valid and reliable central forecast and a good understanding of the uncertainty associated with the forecast. Both of these outcomes require there to be enough model runs (known as iterations) for each simulation to gain confidence in the validity and reliability. However, the model is computationally intensive, so more iterations mean each simulation takes longer to run.

We tested the model by varying the number of iterations, capturing both the model outputs (from five model runs per number of iterations tested, each with a different random seed) and run time. We established that at least 500 iterations would be needed to have good confidence in the validity and reliability of the central forecast and a good understanding of the forecast uncertainty. We also concluded that further increases in iterations could add substantially to run time without significant reliability benefits. The reliability of comparisons *between* different scenarios is supported by minimising randomness and noise between different scenarios. This is achieved by always using the same set of random seeds throughout sets of model runs.

References

Allen, R. *et al.* (2016) *The longer-term costs and benefits of different initial teacher training routes*. Available at: https://ifs.org.uk/sites/default/files/output_url_files/R118.pdf (Accessed: 25 June 2024).

CFE Research and FFT Education Datalab (2023) *Evaluation of the phased maths bursaries pilot: final report*. Available at: https://www.gov.uk/government/publications/evaluation-of-the-phased-maths-bursaries-pilot-final-report (Accessed: 18 November 2024).

Department for Education (2017) 'Teachers' student loan reimbursement: guidance for teachers and schools'. Available at: https://dera.ioe.ac.uk/id/eprint/30399/1/Teachers%27%20student%20loan%20reimbursement_%20guidance%20for%20teachers%20and%20schools%20-%20GOV.UK.pdf (Accessed: 13 June 2025).

DfE (2019) *Early-career payments for teachers*, *GOV.UK*. Available at: https://www.gov.uk/guidance/early-career-payments-guidance-for-teachers-and-schools (Accessed: 25 June 2024).

DfE (2020) *Evidence to the School Teachers' Review Body (STRB): 2020 pay award for school staff*. Available at: https://www.gov.uk/government/publications/evidence-to-the-strb-2020-pay-award-for-school-staff (Accessed: 21 November 2024).

DfE (2022) *Levelling up premium payments for teachers*. Available at: https://www.gov.uk/guidance/levelling-up-premium-payments-for-teachers (Accessed: 25 June 2024).

DfE (2024) *Targeted retention incentive payments for school teachers*. Available at: https://www.gov.uk/guidance/targeted-retention-incentive-payments-for-school-teachers (Accessed: 21 November 2024).

McLean, D., Tang, S. and Worth, J. (2023) *The impact of training bursaries on teacher recruitment and retention: an evaluation of impact and value for money*. Available at: https://www.nfer.ac.uk/media/bycg5uzk/the_impact_of_training_bursaries_on_teacher_recruitment_and_retention.pdf (Accessed: 21 November 2024).

National Audit Office (2016) *Training new teachers*. Available at: https://www.nao.org.uk/wp-content/uploads/2016/02/Training-new-teachers.pdf (Accessed: 25 June 2024).

Worth, J. and Hollis, E. (2021) 'Do bursaries change who applies to teacher training?', *NFER*, 24 November. Available at: https://www.nfer.ac.uk/blogs/do-bursaries-change-who-applies-to-teacher-training/ (Accessed: 25 June 2024).

Worth, J., Tang, S. and Galvis, M.A. (2022) *Assessing the impact of pay and financial incentives in improving shortage subject teacher supply*. Available at:

https://www.nfer.ac.uk/media/t1opzvjm/assessing_the_impact_of_pay_and_financial_incentives_in_improving_shortage_of_subject_teacher_supply.pdf (Accessed: 21 November 2024).

Evidence for excellence in education

Public

The Mere, Upton Park, Slough, Berks SL1 2DQ
T: +44 (0)1753 574123 • F: +44 (0)1753 691632 • enquiries@nfer.ac.uk

www.nfer.ac.uk

NFER ref. GATC

ISBN. 978-1-916567-47-4