

Greaves, Ellen; Sibieta, Luke

Working Paper

Estimating the effect of teacher pay on pupil attainment using boundary discontinuities

IFS Working Papers, No. W14/03

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Greaves, Ellen; Sibieta, Luke (2014) : Estimating the effect of teacher pay on pupil attainment using boundary discontinuities, IFS Working Papers, No. W14/03, Institute for Fiscal Studies (IFS), London,
<https://doi.org/10.1920/wp.ifs.2014.1403>

This Version is available at:

<https://hdl.handle.net/10419/101331>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Estimating the effect of teacher pay on pupil attainment using boundary discontinuities

IFS Working Paper W14/03

Ellen Greaves
Luke Sibieta

The Institute for Fiscal Studies (IFS) is an independent research institute whose remit is to carry out rigorous economic research into public policy and to disseminate the findings of this research. IFS receives generous support from the Economic and Social Research Council, in particular via the ESRC Centre for the Microeconomic Analysis of Public Policy (CPP). The content of our working papers is the work of their authors and does not necessarily represent the views of IFS research staff or affiliates.

Estimating the effect of teacher pay on pupil attainment using boundary discontinuities*

Ellen Greaves[†] and Luke Sibieta[‡]

March 2014

Abstract

This paper provides causal estimates of the effect of teacher pay on pupil attainment using a sharp geographical discontinuity in teacher salaries. We compare schools in close proximity to a pay zone boundary to estimate the effect of teacher salary differentials on pupil attainment. We find that these differences in salary scales do translate into differences in actual teacher pay levels. However, we find little evidence that higher teacher salary scales increase pupil attainment in national assessments at age 11, and are able to rule out quantitatively small effects of 0.07 and 0.02 standard deviations in English and maths, respectively. These results imply that variations in teacher pay of the magnitude we observe (around 5%) are unlikely to be effective for attracting and retaining higher quality teachers.

1 Introduction

Teacher effectiveness, as measured by their pupils' progress, varies dramatically between teachers and appears to be uncorrelated with many observable characteristics of the teacher, such as their level of education (Rockoff (2004); Rivkin et al. (2005); Aaronson et al. (2007); Slater et al. (2009)). Understanding whether mechanisms such as higher teacher pay can be used to attract more effective teachers is therefore important. Estimating the causal relationship between teacher salaries and teacher quality is challenging, however, as there tends to be limited exogenous variation in teacher salaries within countries: teachers' pay is often set in national (or school district-level) agreements, with any variation in salaries tending to reflect education, experience or the nature of the school or costs in the local area (Hanushek and Rivkin (2006)).

This paper makes use of a sharp geographical discontinuity in teacher salary scales in England to identify exogenous differences in teacher pay in otherwise similar areas. Variation arises as teachers in England are currently paid according to centrally determined salary scales, while teachers in the London area receive higher salaries to compensate them for a higher cost of living (and schools receive extra funding in order to pay for these higher salaries). In total, there are four pay zones (inner London, outer London,

*We are grateful to the ESRC for financial support as part of the small grants programme (grant reference ES/J006076/1) as well as the ESRC Centre for the Microeconomic Analysis of Public Policy at IFS (grant number RES-544-28-0001). We thank Rebecca Allen, Steve Gibbons, Carol Propper, Imran Rasul and seminar participants at the Institute for Fiscal Studies, Institute for Education and Public Economics UK 2013 for comments and feedback. The provision of data from the Department for Education is much appreciated. All errors remain our own.

[†]Institute for Fiscal Studies, 7 Ridgmount Street, London. WC1E 7AE, ellen_g@ifs.org.uk.

[‡]Institute for Fiscal Studies and University College London, 7 Ridgmount Street, London. WC1E 7AE, luke_s@ifs.org.uk.

fringe London and rest of England and Wales) and thus three boundaries we could use for our analysis. We focus on just one boundary (fringe London) as the data shows that observables are poorly balanced across the other two boundaries. Teachers at schools close to the fringe London boundary may be on similar points on the pay scale and live in similar neighbourhoods, but will be paid different amounts. We exploit this sharp discontinuity using rich administrative data on pupils, schools, and precise distance to the boundary to examine the effect of teacher salary levels on pupil attainment.

Identification of the causal impact of teacher pay on pupil attainment relies on the similarity of neighbourhoods across the boundary. This identifying assumption is convincing as almost all observable characteristics are balanced. We find little evidence that higher teacher salaries have a positive impact on pupil attainment. Due to the relative precision of our estimates, we are able to rule out quantitatively small positive estimates of 0.07 and 0.02 standard deviations in English and maths at age 11, respectively. This suggests that the sorting of teachers across schools by their quality is not particularly sensitive to salary differentials. This could occur if the decisions of high-quality teachers are not sensitive to salary differentials and/or if schools find it difficult to observe potential teacher quality amongst applicants for a post. The implication of this finding is that offering higher salaries of the magnitude we observe (around 5%) is not likely to be an effective way for schools to improve teacher quality. The structure of rewards (e.g. performance-related pay) or improving the flow of information to schools about the potential quality of teachers could be more effective mechanisms for improving teacher quality.

Our paper is related to a number of literatures: the effect of pay levels on teacher quality; the effect of pay levels on teacher employment decisions; and the effect of school resources.

Hanushek and Rivkin (2006) argue that although a large number of studies attempt to estimate the causal impact of teacher salaries on pupil attainment, most are unable to account for the dependence of teacher salaries on teacher characteristics and/or compensating differentials, or are unable to control for other crucial factors such as past pupil attainment. Amongst the studies they deem to be of high-quality, about 80% show statistically insignificant effects. One study worth highlighting is Loeb and Page (2000) who argue that past estimates of the effect of teacher salaries have been downward biased by a failure to account for differences in non-pecuniary rewards across areas. After accounting for time-constant differences in non-pecuniary rewards, they find that a 10% increase in teacher wages reduces high school drop out rates by 3-4%. Dolton and Marcenaro-Gutierrez (2011) use cross-country comparisons to find a link between teacher salaries and pupil attainment. In a recent study for England, Britton and Propper (2014) exploit collective bargaining for teachers to show that a 10% increase in outside wages for teachers reduces student test scores at age 16 by about 2%, or about 0.1 standard deviations. Taken together, our respective estimates suggest that the effect of pay on pupil attainment is either zero or relatively small, and that this can be found by looking at inside pay (as in our paper) or outside wages (as in Britton and Propper). However, the difference in the point estimates, and particularly those for maths in our paper, provides some suggestion that initial occupational choices might be slightly more sensitive to pay differentials than the sorting of existing teachers.

A related literature considers the effect of relative wages on teacher choices such as entry, duration of teaching and mobility (e.g. Chevalier et al. (2007); Bradley et al. (2012); Dolton (1990); Dolton and van der Klaauw (1995); Dolton and van der Klaauw (1999); Gilpin (2011); Imazeki (2005); Murnane and Olsen (1989); Stinebrickner (1998)). However, Hanushek et al (2004) and Bonesronning et al. (2005) suggest that the effect of student characteristics is much stronger than wages. The results of Hanushek et al (2004) imply

very large compensating differentials, for example a female teacher with 3-5 years experience would require a salary increase of 43% to neutralise the perceived benefits of teaching at a suburban school compared with an urban one. Building on this literature, our paper confirms that the quality of teachers (as measured by pupil attainment) across schools is not particularly sensitive to wage differentials, suggesting that student characteristics and non-pecuniary rewards may be more effective in attracting high quality teachers.

Finally, our paper is related to the literature on the overall effects of school resources. Hanushek (2003) reviews this literature, arguing that there is no strong or consistent relationship between school resources and student achievement. However, some recent papers for England do find small, positive effects of school resources (Holmlund et al. (2010), Machin et al. (2010)). Gibbons et al (2012) use idiosyncratic variation in funding differences across local authorities to estimate a positive effect of resources on attainment. The difference in funding across local authority boundaries varies considerably. However, their results imply that an additional £400 per student per year (a 12.3% increase relative to the mean) could raise achievement by around 0.1 standard deviations. The identification strategy used by Gibbons et al is similar to that of our paper, as funding levels vary across some local authorities in part due to the differences in teacher salary scales¹. Given that our results indicate few differences in pupil attainment as a result of salary scale differences, the results of Gibbons et al are likely to be driven by variation in the quantity of resources.

This paper proceeds as follows: Section 2 describes the key institutional details relating to schools and teachers; Section 3 describes our theoretical framework. Section 4 describes our identification strategy. Section 5 describes our data sources. Section 6 presents some summary statistics and initial graphical results. Section 7 shows our main empirical results. Section 8 concludes.

2 Background: schools and teachers in England

The majority of pupils in England attend primary schools from ages 4-11 before attending secondary school after age 11. Parents can express a preference for which primary school their child attends². We focus on primary schools in our analysis as all pupils in state-funded schools in England must sit external assessments at the end of primary school in Maths and English, known as Key Stage 2 (KS2) tests. Primary schools also tend to be smaller than secondary schools, meaning that there are more close to the boundaries, which aids the precision of our results. Pupils in private or independent schools do not have to sit these tests, but account for only about 6% of each cohort of pupils.

Teachers are not assigned to schools by central or local government. Individual schools post vacancies and teachers apply for these positions. On the basis of these applications, individual schools choose which teacher to hire for a vacant position. If a post is not filled then schools may employ supply or temporary teachers.

The pay and conditions for teachers in England and Wales are regulated by the School Teachers' Pay and Conditions Document (STPCD). During the period covered by our data, the national salary scale had seven different points (M1-4, U1-3)³. In principle, schools have discretion about how quickly their

¹The empirical analysis of Gibbons et al (2012) also considers local authority boundaries where there is no variation in pay-levels.

²Expressing a preference does not guarantee entry to a school: oversubscribed schools use a combination of different rules to allocate places, with the most widely used being whether children have siblings at the school and distance to school.

³M2-3 and U2 have now been removed such that there are effectively two portions of the pay scale, each with minimum and maximum values

teachers move up the pay scale. In practice, during our period of interest, position on the pay scale is largely determined by years of experience. In addition, schools can choose to use some additional payments to pay teachers above the salary scales if they wish⁴, but these flexibilities are relatively under-used by schools.

In order to reflect the higher cost of living in or near London, the level of these pay scales varies across the London area. Specifically, there are four different pay zones: inner London (highest pay zone); outer London; the fringe area of London; and, the rest of England and Wales (lowest pay zone). These four different pay zones are mapped in Figure 1, which also shows the distribution of schools within 2km of each boundary. These boundaries largely coincide with administrative boundaries between Local Authorities⁵. We focus on the schools within 2km of the boundary between fringe London and the rest of England and Wales as the identifying assumption of continuity across the boundary is unlikely to hold for the inner and outer London boundaries. Table 1 shows that under this distance restriction there are 120 schools on the high pay side and 136 schools on the low pay side of the fringe London boundary.

Our empirical strategy relies on the differences in pay scales across *neighbouring* pay zones, shown in Figure 2. By and large, the differential tends to be higher towards the bottom of the pay scale across each of the three boundaries. At the fringe boundary, the difference ranges from 3 to 5% (though the absolute difference is relatively constant at just over £1,000). Pay scales in the fringe area and the rest of England and Wales have grown at exactly the same rate over the period covered by our data (2004-05 to 2010-11) at about 13% in total, so we are not able to exploit any differences in salary scales across time.

Schools in England choose spending levels on individual items subject to national pay and condition for teachers and other rules (e.g. maximum class sizes for under 7s). Over the period of study individual schools were provided with annual budgets from local authorities, each of which has its own school funding formula: each formula allocates a certain amount of funding per pupil, but can vary the basic amount, additional funding for disadvantaged pupils and other factors. In contrast to school districts in the US, local authorities do not raise most of the funds provided to schools through local taxation. Instead, local authorities receive grants from central government, all of which must be spent on schools. These grants are meant to reflect levels of educational need (e.g. levels of social deprivation) adjusted for the local cost of inputs. The adjustment factor is known as the Area Cost Adjustment and is meant to be proportional to differences in teacher salary scales, the costs of other staff and other non-staff costs like energy or books. Therefore, schools in pay zones with higher teacher pay scales are meant to receive a sufficient uplift in funding to allow them to pay higher salaries. However, there are anomalies in the system such that schools in some local authorities are expected to pay teachers more, but not receive a funding uplift⁶. Gibbons et al (2012) exploit such anomalies and other differences in school funding across local authorities to estimate the impact of expenditure per pupil on pupil attainment, finding positive effects of extra resources. The difference in funding across local authority boundaries varies considerably, with the standard deviation of differences in expenditure per pupil across boundaries being around £350 per pupil or about 9% of average per pupil funding. However, their results imply that an additional £400 per student per year (a 12.3%

⁴Additional payments include recruitment and retention payments, teaching and learning responsibility payments and payments for teachers working with children with special educational needs.

⁵Local Authorities have some autonomy, and so there may be variation in council tax and housing policy across these boundaries that could lead to differences in pupil composition on either side of the boundary.

⁶The North London local authority Haringey, for example, receives an area cost adjustment of outer London but is in the inner London teacher pay zone, leading to lower per pupil spending: <http://www.haringey.gov.uk/http://www.haringey.gov.uk/council-fights-for-a-fairer-deal-for-its-schools.htm>

increase relative to the mean) could raise achievement by around 0.1 standard deviations.

The fact that local authorities are responsible for school funding has two further implications. First, the boundary between the fringe area and the rest of England and Wales cuts through some local authorities, some schools in the fringe area are therefore partly dependent on local authorities recognising differences in salary scales in their funding formula. Secondly, local authorities could allocate more funding to schools close to the low-pay side of teacher pay zone boundaries in order to compete with schools nearby on the other side of the boundary. We therefore concentrate on boundaries where there are differences in teacher salary scales and measure the extent to which the funding differential across boundaries is sufficient to compensate schools for the higher salary scales.

3 Theoretical predictions

This section outlines the theoretical predictions for considering how differences in teacher salary scales imposed across boundaries may affect pupil attainment, as a result of the actions of multiple actors. This informs our interpretation of the estimated differences in attainment across the fringe London boundary.

We assume that pupil attainment is a function of observable pupil attributes (including family background), area attributes and teacher effectiveness.

$$Y_i = f(X_{pi}, X_{ai}, S_i T_i, \varepsilon_i) \quad (1)$$

Where X_{pi} represents pupil characteristics, such as their socio-economic background, X_{ai} represents area characteristics, such as local deprivation, S_i represents school attributes such as resources, T_i represents teacher effectiveness and ε_i represents all unobservable influences pupil attainment such as their ability.

Interest lies in the direct effect of salary levels on T_i on Y_i . We cannot measure T_i directly following the standard methodology in the literature as it is not currently possible to link individual teachers to pupils in data for England⁷. Instead, we rely on across school differences in mean pupil outcomes and teacher salary scales (induced by the teacher pay boundary, which we show leads to an increase in teacher pay on the high pay side of the boundary given the high regulation of teacher pay in England) to estimate the causal effect of teacher salary levels on pupil attainment. In order to interpret such an effect as working via changes in teacher quality, we must argue that all other possible influences on pupil attainment are held constant, which depends on the actions of other agents. We now discuss in turn the theoretical predictions for the actions of local authorities, schools and teachers, discussing how each agent on the high and low pay side of boundaries may respond.

3.1 Local Authorities

Local authorities on the high-pay side of pay zone boundaries should receive additional funds from central government (through the 'Area Cost Adjustment') to allow schools in their area to pay the higher salary scales. The teacher pay bill is typically a large proportion of gross total expenditure: around 50% for primary schools in England in 2010-11. This means that local authorities in the high pay area need to provide around

⁷Given that within school variation in salary levels is likely to reflect teacher characteristics such as experience, it would be difficult to isolate exogenous variation in within school salary levels in any case.

50% of the difference in the teacher salary scales in order to allow their schools to afford the same bundle of teacher versus non-teacher resources. We will test to what extent the funding differential matches this figure. If the funding differential provided by local authorities is more (less) than about 50% of the difference in teacher salary scales, schools will be over (under) compensated for the difference in salary scales.

Another potential response is for local authorities to allocate more funds to schools on the low-pay side of the boundary to offset any perceived disadvantage. We test for this by looking at the actual differences in resources and funding in schools either side of the pay boundaries, compared with the actual differences in salary scales.

3.2 Schools

The response of schools will partly be determined by the response of local authorities. If schools are over-compensated for the salary differential they may employ a greater quantity of resources, or fewer if they are under-compensated. However, even if they are perfectly compensated there is still potential for a negative substitution effect as teachers are more expensive relative to other educational inputs for schools on the high-pay side of each boundary (assuming that the prices of other goods, such as books, are constant across the teacher pay boundary). All else being equal, this will lead schools to spend a lower proportion of their income on teachers. Non-proportional differences in the salary scales could affect the mix of teachers by experience. In practice, however, schools are bound to extent by national government regulations (e.g. statutory limits on class sizes for pupils aged 7 and under).

We test for the presence of these substitution effects by looking at the quantity and mix of resources chosen by schools in terms of the pupil-teacher ratio, the pupil-assistant ratio and the balance between teachers and assistants.

Individual schools do have some flexibility in setting teacher salaries and therefore one further potential response for schools on the low pay side of the boundary is to change their pay policies. For instance, schools on the low-pay side of the boundary could pay extra allowances to teachers or accelerate teachers through the salary scales to compensate for the salary differential with nearby schools, effectively smoothing out the salary differential, which we test for.

3.3 Teachers

Differences in teacher salary scales either side of each pay boundary could influence quality of teachers, and thereby pupil attainment. This could occur through two main potential channels: selection-effects; and, efficiency-wage concerns.

First, to illustrate the selection effects let us consider two schools that are observationally equivalent except for their level of teacher pay, where school H is in the high-pay region and school L is in a neighbouring lower pay region. We assume a set of potential teachers who differ in terms of their outside option (u_j) and their level of intrinsic motivation (η_j) as per the motivated agents of Besley and Ghatak (2005) and Bó et al. (2012). In the context of a simple Roy model, potential teachers will then apply to school i if the following individual rationality condition holds, where w_i is the salary offered at school i .

$$w_i + \eta_j > u_j \tag{2}$$

Assuming that an individual's utility from working at school i follows this form, there will be three sets of potential teachers. The first set will apply to both sets of schools as they are either highly motivated and/or have worse outside options. The second set will contain teachers who only apply to the high pay school. The third set will be individuals who apply to neither school, either because they have better outside options or gain little direct utility from being a teacher.

The difference in the average quality of teachers employed at the two schools depends on the observability of teacher quality amongst the two pools of applicants. If schools can easily observe the potential quality of applicants, then offering higher pay can't lead to a decline in average teacher quality. However, if teacher quality is unobservable, then quality could be lower in school H in the case where teacher quality is strongly linked to intrinsic motivation and such motivation is relatively uncorrelated with outside options (as is pointed out by Bó et al. (2012) for the case of public servants more generally). For example, consider the extreme case where motivation and outside options are completely uncorrelated. In this case, school H will have a larger pool of applicants with higher outside options and lower levels of motivation, on average. If teacher quality is unobservable and more strongly linked to motivation, then the larger pool of applicants increases the chances of selecting a lower quality teacher. The overall effect on teacher quality via the selection channel is thus ambiguous. The direction of the effect depends on the observability of potential quality to schools, and the relative importance of intrinsic motivation in shaping teacher quality. The fact that quality could decline in such situation has been shown by Delfgaauw and Dur (2007) and has been applied to the market for nurses by Heyes (2005).

Second, a difference in wages across a teacher pay boundary could increase the quality of teaching via efficiency wage effects (Shapiro and Stiglitz (1984)). Although separation rates for unsatisfactory performance are very low, teachers could still be allocated unpleasant tasks if low effort levels are observed.

Therefore, the overall effect of the salary differential on average teacher quality is ambiguous, depending on the size and direction of the selection effect and the relative effect of efficiency wages.

4 Empirical Methodology

Our goal is to estimate the causal effect of higher teacher salaries on pupil attainment. As indicated by Hanushek and Rivkin (2006) and Loeb and Page (2000), it can be difficult to find exogenous variation in teacher salaries. Firstly, differences in teacher salaries across areas could reflect compensating variations due to differences in the cost of living or non-pecuniary benefits (Rosen (1987)). Secondly, within-school variation in teacher salaries could just reflect differences in teacher characteristics, such as progression through the salary scales due to experience.

Our identification strategy relies on a sharp regression discontinuity design, where we rely on continuity assumptions at the boundary between treatment and control schools. In particular, we compare schools either side and within close proximity of the fringe London teacher pay zone boundary, where the difference in attainment is our estimate of the causal effect of differences in teacher salary levels on pupil attainment. For this strategy to work, we require schools to well balanced in areas in close proximity to the boundary. Below we outline our identification strategy and empirical methodology.

4.1 Identification strategy

For ease of exposition, let us assume that the production function in equation 1 is additive and separable. If we then take the difference in mean outcomes across schools in the high and low-pay regions, we have the following:

$$\bar{Y}_H - \bar{Y}_L = \beta_p (\bar{X}_{pH} - \bar{X}_{pL}) + \beta_a (\bar{X}_{aH} - \bar{X}_{aL}) + \beta_S (\bar{S}_H - \bar{S}_L) + \beta_T (\bar{T}_H - \bar{T}_L) + (\bar{\varepsilon}_H - \bar{\varepsilon}_L) \quad (3)$$

where \bar{x} denotes the mean of x . By comparing schools in very close proximity to the boundary, we argue that $\bar{X}_{aH} = \bar{X}_{aL}$. For example, the cost of living is likely to be roughly constant across the boundary. Similarly, we expect to find very small differences in observable pupil characteristics ($\bar{X}_{pH} - \bar{X}_{pL}$) and unobservable characteristics ($\bar{\varepsilon}_H - \bar{\varepsilon}_L$). If local authorities properly compensate schools for the difference in salary scales, then we also expect $(\bar{S}_H - \bar{S}_L)$ to be zero. If these assumptions hold, the difference $\bar{Y}_H - \bar{Y}_L$ should therefore represent the difference in teacher effectiveness driven by the exogenous difference in teacher pay across the boundary. This is effectively a restatement of the continuity assumption required for identification in regression discontinuity designs (Blundell and Dias (2009); Lee and Lemieux (2010)).

The only exception to the continuity assumption is that we expect to see a discontinuity in teacher pay levels and cash-terms funding levels: schools on the high-pay side of the boundary should receive higher levels of funding in order to allow them to pay teachers according to the higher salary scales. If these funding levels are overly generous, then schools on the high-pay side of the boundary could have higher levels of real resources or increase the size of the actual salary differential relative to the difference in salary scales. As explained in section 3.1, with no expected differentials in other inputs prices, a funding differential of around 50% of the salary differential would be approximately sufficient to allow schools either side of the pay zone boundaries to purchase the same bundle of resources.

4.2 Empirical Methodology

We compare pupil outcomes in schools either side and within close proximity of the fringe London teacher pay zone boundary. Following Lee and Lemieux (2010), we explicitly check the balance of covariates either side of the boundary, test the sensitivity of results to controlling for covariates and their sensitivity to different definitions of being 'close' to the boundary (equivalent to varying the bandwidth in standard RDD estimates).

Our empirical analysis begins by illustrating the raw differences in pupil outcomes and covariates across the boundary. We present some of the standard RDD graphs recommended by Lee and Lemieux (2010) to illustrate how pupil outcomes change close to the teacher pay boundary. Figures 3 and 4 do not show a strong relationship between distance to the boundary and pupil outcomes (and we are unable to reject the null hypothesis that the coefficient on distance to the boundary is zero). This is to be expected as local labour market conditions are unlikely to vary much within such small distance. For this reason, our main analysis then focuses on a simplified RDD design, where we simply compare the outcomes of schools within close proximity to the boundary.

We focus our analysis on the boundary between fringe London and the rest of England where observable characteristics are well balanced: Table 1 illustrates the balance in observable covariates between

schools either side of this boundary, and the imbalance of pupil characteristics across the inner and outer London boundaries, which limits our confidence in the balance of *unobservable* characteristics across these two boundaries.

Our point estimates are unchanged when we control for small differences in covariates at the fringe boundary, which validates the identification assumption of continuity in covariates across boundaries. The benefit of controlling for covariates is that our estimates become more precise.

We use various methods to control for differences in observable characteristics. First, we use linear regression: assuming that the impact of pupil and area characteristics on pupil attainment is the same on the high and low-pay side of the boundary, Equation 3 can be written:

$$Y_s = \beta_p X_{sp} + \beta_a X_{sa} + \beta_S S_s + \beta_T T_s + \varepsilon_s \quad (4)$$

Where the subscript s represents schools (the average of all pupils i at the school). β_T represents the impact of teacher effectiveness, driven by exogenous differences in teacher pay on the high and low-pay side of the boundary. β_T is consistent if the following conditional independent assumption holds: $E(\varepsilon_s | X_{sp}, X_{sa}, T_s) = 0$. This is more plausible in circumstances where observables are also balanced, but is ultimately untestable. The requirement that there is no difference in conditional unobservable characteristics is common to all approaches we adopt.

A more flexible approach to controlling for observable characteristics is Fully-Interacted Linear Matching (FILM). This specification interacts the treatment indicator T_s with all observable characteristics, X_{sp} and X_{sa} . This allows the treatment effect to vary with observable characteristics.

The third approach is propensity score matching, which is more flexible than FILM and OLS. A suitable comparison group for each school s is constructed by matching those on the high and low-pay side of the boundary. This approach explicitly drops schools with no suitable comparators (those with no common support), while the two methods above extrapolate based on the functional form specified (masking problems of common support). The validity of propensity score matching depends on the conditional independence assumption, where the outcome is uncorrelated with assignment to treatment, conditional on observable characteristics.

For the main results we include in the estimation sample only schools within 2 km of each boundary. As a robustness check, we examine the sensitivity of the results to including schools further away (3 km) and only using schools very close to the boundary (1 km). Distance to the boundary is also included as an observable characteristic to achieve a reasonable balance of schools very close to the boundary, although, as already indicated, distance to the boundary does not seem to have an independent effect on attainment (conditional on being within 2km of the boundary).

5 Data

We link together data from a number of administrative data sets over various years. In particular, we use data from the National Pupil Database and Spring Census from 2006 to 2011, which contain the test results and observable characteristics of every pupil in state-funded schools in England. Our main outcomes are the school-level average points scored in Key Stage 2 Maths and English, standardised at the national level

(externally assessed Science tests were stopped from 2009 onwards). We disregard data for 2010 as a large number of schools boycotted Key Stage 2 national examinations in that year. Our sample consists of schools with non-missing Key Stage 2 results who remain in the sample for all years from 2006 to 2011 (excluding 2010) and who are 'close' to one of the pay boundaries.

We are also able to derive various school level characteristics: number of pupils; proportion of pupil eligible for free school meals (FSM); proportion of pupil with English as an additional language (EAL); proportion of pupils with statements of special educational needs (SEN), with and without statements; and, proportion of pupils from non-white backgrounds. Eligibility for free school meals is a rather coarse measure of deprivation, so we also use other measures of deprivation based on the area in which pupils live: average percentile rank on the Index of Multiple Deprivation, and its various domains (education, employment, income); average percentile rank on the Income Deprivation Affecting Children Index (IDACI).

We use funding and expenditure levels defined in Section 251 outturn data, which reports funding and expenditure levels for each financial year (April to March) for all maintained schools in England. Information on staffing levels is taken from the Local Education Authority School Information Service (LEASIS) and its later replacements.

As Key Stage 2 tests are taken in the summer of each year, we link these results to school characteristics defined in January of the same year (taken from the Spring Census), staffing levels defined for same academic year (LEASIS) and to funding/expenditure levels in the financial year most recently ended (e.g. we link May 2011 test results to funding and expenditure per pupil for the 2010-11 financial year).

For teacher pay levels, we make use of the School Workforce Census that contains the pay and characteristics for all employees in schools across England from 2010 onwards. As this survey is still relatively new, we only make use of data from 2011.

6 Descriptive evidence

Before presenting the detailed empirical results in section 7, we briefly present some summary statistics of the raw differences in outcomes and observable characteristics across each pay zone boundary. As discussed in section 4, we would like to see that the differences in observable characteristics across each boundary are insignificant or small as this makes our identification assumptions somewhat more plausible. We also present some graphical analysis showing outcomes across pay zone boundaries.

6.1 Summary Statistics

Table 1 shows the average characteristics of schools within 2km of each pay zone boundary, and compares the characteristics of schools just inside and outside each boundary. This is done separately for the inner, outer and fringe pay zone boundaries. The first two columns show the average characteristics of schools across England and London for reference. Figures are pooled across all years. Changes over time in covariates and outcomes at each of the boundaries are presented in the online appendix, which we refer to below.

For the differences across the inner London pay zone boundary, funding and expenditure per pupil are significantly higher on the high-pay side of the boundary, with these differences (14%) well in excess of the relative differences across all points in the salary scale (4-11%) and that required to compensate schools for the higher pay for teachers. It may be surprising, therefore, that maths and English test scores are

significantly lower on the high-pay side of this boundary. However, it is also clear that there are large and significant differences in observable characteristics of pupils across the boundary, which may contribute to the differences in attainment. For example, there is a 9 percentage point gap in the proportion of children eligible for FSM and a 10 percentage point gap in the proportion of pupils from non-white backgrounds. These differences are large, representing around half the national standard deviation in each case. The online appendix also shows that the differences in covariates have remained relatively constant over time. However, schools on the high-pay of the inner London boundary have improved significantly over time, particularly in maths. The changing size of the differential in outcomes over time thus suggests that unobservables have also been shifting over time. We therefore make no claim that we can identify the effect of salary differences along this boundary.

There is a similar pattern across the outer London pay boundary, although the differences are slightly less extreme: schools on the high-pay side of the boundary are significantly more deprived than schools just outside the boundary, and have significantly more pupils from non-white backgrounds. However, despite being more deprived, schools on the high-pay side of the boundary have slightly higher test scores, significantly so in maths. We also see large differences in funding (13%) in excess of the salary differential across all points in the scale (7-11%) and that required to compensate schools for the higher pay for teachers. There are also significantly lower pupil:teacher and pupil:assistant ratios on the high-pay side of the boundary. Higher pupil attainment could thus be a function of higher resources or higher levels of teacher pay. Combined with the large differences in observables, it is thus not possible to use this boundary to identify the effect of pay on attainment.

Finally, at the fringe area boundary, it is clear that schools are largely balanced in terms of observable characteristics, though there are some small and significant differences in wider socio-economic measures (the Index of Multiple Deprivation). There are no significant differences in raw funding levels, which may suggest that these schools do not receive extra funding to pay for the higher teacher salary scales. There are also no significant differences in test scores. In principle, this could reflect the lack of any effect of the pay differential, but could also reflect a lack of resources to finance the pay differential. Below, we therefore investigate the differences in funding levels in more detail. We also address this concern directly by showing that there are differences in actual teacher pay levels and that there is no evidence of pay smoothing. Reassuringly, the online appendix figures show that observable characteristics and pupil outcomes across the fringe boundary are balanced in all years.

In summary, the inner and outer London pay boundaries are poorly balanced in observable characteristics and exhibit different trends in pupil attainment over time, making it unlikely that unobservable characteristics are balanced. For outer London, we see slightly higher outcomes on the high-pay side of the boundary, which could result from a teacher pay or resources effect (given the overly-generous funding), or from an imbalance in unobservable characteristics. There is a relatively good balance in observable characteristics along the fringe boundary, making it an ideal candidate to test the effect of salary differentials. From now on, we therefore focus attention on the fringe boundary, but make reference to the equivalent results for inner and outer London, which are contained in the online appendix.

6.2 Graphical Analysis

It is also helpful to show how outcomes vary depending on distance from the fringe boundary. This illustrates the data and methods we are using, and informs the direction of empirical analysis in the next section. To do so, we use graphs recommended by Lee and Lemieux (2010) for RDD methods. In particular, we show the local averages for schools in bins of size 200m in terms of distance either side of the fringe boundary (black dots) up to 3 km from the boundary, as well as estimates of the relationship between distance to the boundary and each outcome based on a linear specification (dashed line) and a 7th order polynomial (solid line), each with a break at the discontinuity. We show this for English (3) and maths (4), with data pooled across years.

In both cases, there is no clear or consistent relationship between test scores and distance from the pay boundary (at least within 3km either side of the pay boundaries). The relationship between distance to the boundary and attainment is best described by a flat line with noise, with the high order polynomial oscillating around the linear estimates. Indeed, in a linear regression we are unable to reject the null hypothesis that the slope coefficients on distance are zero either side of the boundary. There is also little evidence to suggest a positive jump at the pay boundary. As a result, our main empirical analysis represents a simplified version of RDD where we compare schools within 2 km of the boundary, arguing that such schools are likely to be very similar on observables (which we have already shown) and on unobservables (which we cannot test, but argue is likely to hold in this context).

7 Empirical Analysis

This section presents our main empirical results for the effects of discontinuities in teacher pay on pupil attainment. We focus on the fringe boundary, but also make reference to the equivalent results for the other two boundaries (which are contained in the online appendix). We tie the results back to our theoretical predictions. In particular, we start by looking at the potential responses of local authorities by looking at differences in funding, comparing these with the differences in the salary scales to assess the extent to which schools on either of a pay boundary can afford the same bundles of resources. We then look at potential schools responses in terms of the actual resources employed (in terms of staff to pupil ratios) and whether the differences in salary scales translate into actual teacher pay levels (and other aspects of their remuneration). This is important to address two concerns: first, whether schools respond to the discontinuity in teacher pay by changing their pay policies, e.g. paying more to teachers on the low-pay side of the boundary. Second, whether schools on the high-pay side of the fringe boundary actually pay their teachers more.

Having examined differences in resources and actual teacher pay levels, we then address our main question: whether discontinuities in teacher salaries translate into differences in pupil attainment.

For each outcome, we estimate the difference in outcomes across the fringe boundary using four different methodologies (Raw, OLS, FILM, Matching) as described in section 4. We also estimate the differences for each individual year separately, which are presented in the online appendix. Here, we just show the results for the latest year of data (2011). In the final section, we vary the distance to the boundary to demonstrate that our results are not dependent on the precise measure of distance chosen.

7.1 Resources

The first panel of Table 2 shows the differences in differences in (log) funding per pupil, pupil:teacher ratios and pupil:assistant ratios at the fringe boundary. As one can see, there remains only a small difference in funding levels at schools either side of the boundary. This remains true across all specifications. Given a difference in salary scales of around 4-5%, we expect schools on the high-pay side of the boundary would need about 2-3% extra funding in order to afford the same bundle of resources. Although we cannot reject that this is indeed the case, the actual differences suggest that schools on the high-pay side of the boundary might be unable to pay their teachers more or might employ fewer resources in total. On the first point, schools are not able to pay less than statutory pay levels (though they may be able to move teachers more slowly through the pay scales). On the second point, schools on the high-pay side of the boundary actually seem to have more resources, at least in terms of having lower pupil:teacher and pupil:assistant ratios (though the differences are insignificant).

The online appendix shows the results for inner and outer London. For inner London, this shows that the difference in funding per pupil is generous compared with the salary scales, but that there is little difference in actual resources employed. At the outer London boundary, we see differences in funding per pupil of around 11-13% in 2011, which are much larger than those in the salary scales. These translate into substantial differences in actual resources: the pupil:teacher ratio is around 0.9 lower on the high-pay side of this boundary, about one third of the standard deviation in the pupil:teacher ratio in London as a whole

7.2 Teacher Pay

It is important to demonstrate that the discontinuities in pay scales translate into differences in actual levels of teacher pay, particularly in light of the lack of difference in funding levels at the Fringe boundary. Therefore, the middle panel of Table 2 shows the difference in the total pay of individual classroom teachers at each of the the fringe boundary, before and after controlling for school characteristics.

We see that the raw difference in actual teacher pay levels is estimated at just over £900. This almost perfectly in line with the difference in salary scales at this boundary (ranging around £1,000 to £1,100). These differences are smaller when we control for observable characteristics. However, there are no differences across boundaries in terms of the actual margins available to schools when setting teacher pay. For example, there are no significant differences in average spine point of teachers at the Fringe boundary (indeed, the point estimates suggest teachers on the high pay side of the boundary are on slightly higher spine points, not lower ones) and there are no significant differences in the proportion of teachers with additional responsibility payments (one flexibility that could be used by schools to smooth the difference in salary scales). Therefore, the difference in salary scales at at the fringe boundary does seem to translate into actual differences in teacher pay, with little evidence of pay smoothing.

As shown in the online appendix, at the inner and outer London boundary, the actual difference in pay is slightly less than that in the salary scales, suggesting some pay smoothing.

7.3 Attainment

The bottom panel of Table then shows the estimates for the differences in average pupil attainment in maths and English.

This shows little evidence of any positive impact of the higher teacher pay area around the fringe boundary: across all specifications, there is no evidence of a positive effect on pupil attainment. For English, the point estimates are around 0 to 0.01, which are estimated very precisely in the case of FILM (standard error of 0.03). The point estimates for maths are generally very small and negative across all specifications. The 95% confidence intervals implied by these estimates mean that we are able to rule out quantitatively small effects of 0.07 and 0.02 standard deviations in English and maths. As shown in the online appendix, these differences are very stable over time for both maths and English.

Combined with our previous results, these findings suggest there is no evidence that higher teacher salaries improves teacher effectiveness in terms of pupil attainment.

At the inner London boundary (as shown in the online appendix), the differential in attainment is small and insignificant once we control for observables. However, it is hard to interpret this as evidence that higher teacher has no impact on pupil attainment: the balance of unobservable characteristics is questionable given the lack of balance in observable characteristics, and is likely to be changing over time given the dramatic differences in attainment over time. At the outer London boundary, there is a significant and positive impact of being on the high-pay side of the boundary for schools around the outer London boundary once we control for observable characteristics in OLS and FILM. However, there was poor balance in observables at this boundary, as well as a large difference in resources, making it difficult to interpret this as the causal effect of higher salary scales.

7.4 Robustness checks

Table 3 and 4 show our estimates of the effect of the pay differential at the Fringe Area boundary for different measures of being close to the boundary and across all years for maths and English, respectively. In particular, we show the effect of the pay differential for schools within 1, 2 and 3 km of the Fringe Area boundary. This shows that the qualitative conclusions are largely unchanged if we just look at schools within 1 km of the pay boundary or expand our analysis to also look at schools within 3 km, and the differences are all stable over time.

7.5 Discussion of results

The main findings from our analysis is the lack of evidence to suggest that higher teacher salaries improves teacher effectiveness in terms of pupil attainment. How does this compare with other estimates in the literature? For the US, Loeb and Page (2000) find that a 10% increase in teacher wages reduces high school drop out rates by 3-4%, a relatively large effect considering that the state-level standard deviation in drop out rates ranged from 6.1% in 1959 to 2.5% in 1989 at the end of the sample period. In comparison, our results are small and insignificant. For the UK, Britton and Propper (2014) find that a 10% increase in the outside wage depresses test scores at age 16 by 2%, or about 0.1 of the standard deviation in test scores. Our results imply that a 5% difference in salary scales is consistent with an increase in test scores up to 0.02 standard deviations for maths and 0.07 standard deviations for English. Our results for English thus seem consistent with those found by Britton and Propper, though our results for maths are lower.

There are three reasons that would lead us to expect to see differences between our papers. First, there is a difference in focus, we examine primary schools, whilst Britton and Propper look at secondary

schools. It may be that primary school teachers behave differently to secondary school teachers, or there may be differences in their outside options. Indeed, secondary school teachers more likely to have completed a degree in a relevant subject, followed by a one-year conversion course, and may thus have better outside options than primary school teachers, who are more likely to have completed a three-year course in education. Second, we use different identification strategies. We would argue that our methodology based on sharp boundary discontinuities provides causal estimates of the effect of the differences in teacher salaries. Moreover, for the fringe boundary, our estimates are relatively precise, we show that covariates are well-balanced and there is no evidence of pay smoothing. Britton and Propper rely on differences in outside wages across regions over time. They therefore require that differential trends in outside wages across regions are not associated with any unobservable heterogeneity. They also assume that inside wages do not respond to these outside wages as there are national pay structures, which is plausible given the rare use of pay freedoms over this period.

Third, and perhaps most importantly, there is a difference in the sorting mechanisms that could be driving the results. In relying on differences in outside wages, Britton and Propper’s results are likely to be driven by sector or occupational choices. For instance, their results suggest that when outside wages are high, individuals going into teaching are likely to be of lower average ability, which is consistent with previous evidence on how relative wages affects the sorting of teachers by ability (Hoxby and Leigh (2005); Nickell and Quintini (2002); Bacolod (2007)). However, their results are unlikely to be driven by existing teachers sorting across schools within close proximity as these choices are less likely to be affected by the value of the local outside option. In contrast, our results are likely to be primarily driven by teachers sorting across schools in a local area, with the conclusion that local salary differentials have little effect of the sorting of teachers in terms of their overall effectiveness, consistent with other evidence suggesting that existing teacher decisions are more sensitive to pupil or school characteristics (e.g. Hanushek et al. (2004)).

The comparison of the results is thus informative. Our respective estimates suggest that the effect of pay on pupil attainment is either zero or relatively small, and that this can be found by looking at inside pay (as in our paper) or outside wages (as in Britton and Propper). However, the difference in the point estimates, and particularly those for maths, provides some suggestion that initial occupational choices might be slightly more sensitive to pay differentials than the sorting of existing teachers.

8 Conclusion

In this paper, we sought to use sharp geographical discontinuities in teacher salary structures in England to identify the effect of teacher pay levels on pupil attainment. Identification of the causal impact relies on a standard continuity assumption. There is only one boundary where this identifying assumption is convincing, however, in that almost all observable characteristics are balanced. At this boundary, between the fringe area and the rest of England, no significant impact of teacher pay is evident, and there is no evidence of differences in resources or pay smoothing.

Our conclusions on the effect of teacher pay suggest that the sorting of teachers across schools by effectiveness is not particularly sensitive to universal salary differentials. This could occur if the decisions of high-quality teachers are not sensitive to salary differentials and/or if schools find it difficult to observe potential teacher quality amongst applicants. This implies that offering higher universal salaries is unlikely

to be an effective way for schools to improve teacher quality. More effective strategies could include changing the structure of rewards (e.g. performance-related pay), improving information provided to schools about applicants or greater firing of existing teachers deemed to be low quality. There is a growing literature on the structure of teacher rewards, particularly the effectiveness or otherwise of performance-related pay. Most recent estimates find a positive effect of performance-related pay on pupil achievement (Atkinson et al. (2009); Figlio and Kenny (2007); Glewwe et al. (2010); Ladd (1999); Lavy (2002); Lavy (2009); Muralidharan and Sundararaman (2011); Podgursky and Springer (2007)), though these findings are not universal (Eberts et al. (2002)). Set against our findings, this implies that varying the structure of rewards may be a more effective policy lever than varying the level of base rewards.

In future work, it will be important to understand the reasons for observing no effect of pay on pupil attainment. A limited effect of pay could reflect a situation where potential teacher quality is not observable amongst a pool of applicants. Investigating this explanation would require collecting data on the characteristics of the pool of applicants for individual teacher positions either side of the pay zone boundaries. This would allow us to see if the pay differentials have an impact on the size and quality of applicant pools. If there is no effect, then the more likely explanation would be that teachers do not weight pecuniary rewards particularly highly when making employment decisions.

9 Bibliography

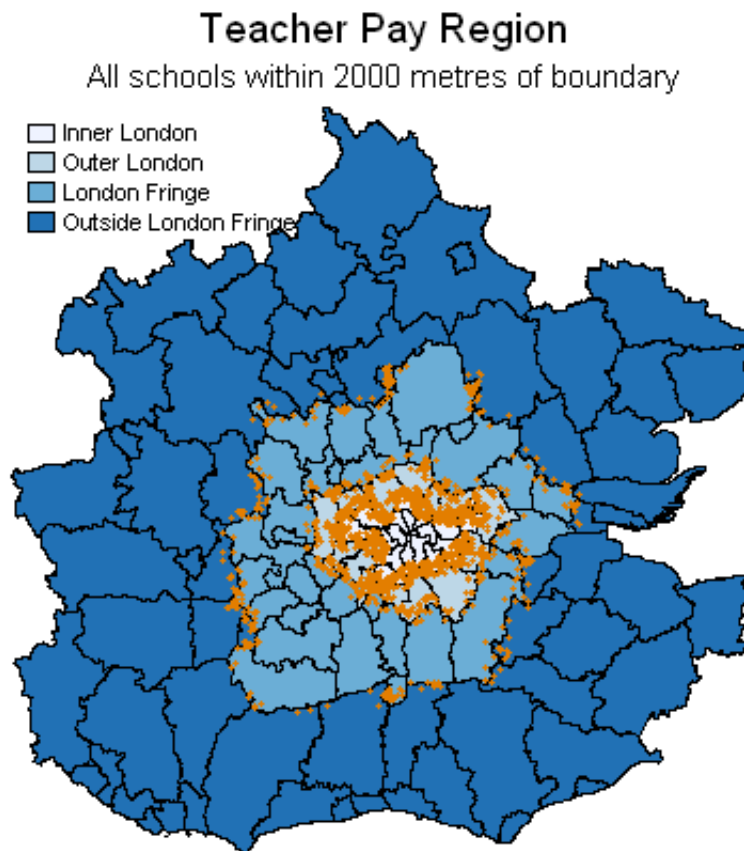
- Aaronson, D., L. Barrow, and W. Sander (2007): “Teachers and Student Achievement in the Chicago Public High Schools,” *Journal of Labor Economics*, vol. 25(1), 95–135.
- Atkinson, A., S. Burgess, B. Croxson, P. Gregg, C. Propper, H. Slater, and D. Wilson (2009): “Evaluating the impact of performance related pay for teachers in England,” *Labour Economics*, vol. 16, 251–261.
- Bó, E. D., F. Finan, and M. Rossi (2012): “Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service,” NBER Working Papers 18156, National Bureau of Economic Research, Inc.
- Bacolod, M. (2007): “Do Alternative Opportunities Matter? The Role of Female Labor Markets in the Decline of Teacher Quality,” *Review of Eco*, vol. 89(4), 737–751.
- Besley, T. and M. Ghatak (2005): “Competition and Incentives with Motivated Agents,” *American Economic Review*, vol. 95(3), 616–636.
- Blundell, R. and M. C. Dias (2009): “Alternative Approaches to Evaluation in Empirical Microeconomics,” *Journal of Human Resources*, vol. 44(3).
- Bonesronning, H., T. Falch, and B. Strom (2005): “Teacher sorting, teacher quality, and student composition,” *European Economic Review*, vol. 49(2), 457–483.
- Bradley, S., C. Green, and J. Mangan (2012): “The Effect of Relative Wages and External Shocks on Public Sector Turnover,” *The Economic Record*, vol. 88(281), 243–261.
- Chevalier, A., P. Dolton, and S. McIntosh (2007): “Recruiting and Retaining Teachers in the UK: An Analysis of Graduate Occupation Choice from the 1960s to the 1990s,” *Economica*, vol. 74(293), 69–96.
- Delfgaauw, J. and R. Dur (2007): “Signaling and screening of workers’ motivation,” *Journal of Economic Behavior & Organization*, vol. 62(4), 605–624.
- Dolton, P. and O. D. Marcenaro-Gutierrez (2011): “If you pay peanuts do you get monkeys? A cross-country analysis of teacher pay and pupil performance,” *Economic Policy*, vol. 26(65), 5–55.
- Dolton, P. and W. van der Klaauw (1999): “The Turnover of Teachers: A Competing Risks Explanation,” *The Review of Economics and Statistics*, vol. 81(3), 543–550.
- Dolton, P. J. (1990): “The Economics of UK Teacher Supply: The Graduate’s Decision,” *Economic Journal*, vol. 100(400), 91–104.
- Dolton, P. J. and W. van der Klaauw (1995): “Leaving Teaching in the UK: A Duration Analysis,” *Economic Journal*, vol. 105(429), 431–44.
- Eberts, R., K. Hollenbeck, and J. Stone (2002): “Teacher Performance Incentives and Student Outcomes,” *Journal of Human Resources*, vol. 37(4), 913–927.
- Figlio, D. N. and L. W. Kenny (2007): “Individual teacher incentives and student performance,” *Journal of Public Economics*, vol. 91(5-6), 901–914.

- Gilpin, G. A. (2011): "Reevaluating the effect of non-teaching wages on teacher attrition," *Economics of Education Review*, vol. 30(4), 598–616.
- Glewwe, P., I. Nauman, and M. Kremer. (2010): "Teacher Incentives," *American Economic Journal: Applied Economics*, vol. 2, 205–27.
- Hanushek, E. A. (2003): "The Failure of Input-Based Schooling Policies," *Economic Journal*, vol. 113(485), F64–F98.
- Hanushek, E. A., E. F. Kain, and S. G. Rivkin (2004): "Why Public Schools Lose Teachers," *Journal of Human Resources*, vol. 39(2).
- Hanushek, E. A. and S. G. Rivkin (2006): *Teacher Quality, Handbook of the Economics of Education*, vol. 2, chapter 18 (Elsevier), pp. 1051–1078.
- Heyes, A. (2005): "The economics of vocation or 'why is a badly paid nurse a good nurse'?" *Journal of Health Economics*, vol. 24(3), 561–569.
- Holmlund, H., S. McNally, and M. Viarengo (2010): "Does money matter for schools?" *Economics of Education Review*, vol. 29(6), 1154–1164.
- Hoxby, C. and A. Leigh (2005): "Pulled Away or Pushed Out? Explaining the Decline of Teacher Aptitude in the United States," *American Economic Review*, vol. 94, 236–240.
- Imazeki, J. (2005): "Teacher salaries and teacher attrition," *Economics of Education Review*, vol. 24(4), 431–449.
- Ladd, H. (1999): "The Dallas school accountability and incentive program: an evaluation of its impacts on student outcomes," *Economics of Education Review*, vol. 18, 1–16.
- Lavy, V. (2002): "Evaluating the Effect of Teachers' Group Performance Incentives on Pupil Achievement," *Journal of Political Economy*, vol. 110(6), 1286–1317.
- Lavy, V. (2009): "Performance Pay and Teachers' Effort, Productivity, and Grading Ethics," *American Economic Review*, vol. 99(5), 1979–2011.
- Lee, D. S. and T. Lemieux (2010): "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, vol. 48(2), 281–355.
- Loeb, S. and M. E. Page (2000): "Examining The Link Between Teacher Wages And Student Outcomes: The Importance Of Alternative Labor Market Opportunities And Non-Pecuniary Variation," *The Review of Economics and Statistics*, vol. 82(3), 393–408.
- Machin, S., S. McNally, and C. Meghir (2010): "Resources and Standards in Urban Schools," *Journal of Human Capital*, vol. 4(4), 365 – 393.
- Muralidharan, K. and V. Sundararaman (2011): "Teacher Performance Pay: Experimental Evidence from India," *Journal of Political Economy*, vol. 119(1), 39 – 77.

- Murnane, R. J. and R. J. Olsen (1989): “The Effects of Salaries and Opportunity Costs on Duration in Teaching: Evidence from Michigan,” *The Review of Economics and Statistics*, vol. 71(2), 347–52.
- Nickell, S. and G. Quintini (2002): “The Consequences of The Decline in Public Sector Pay in Britain: A Little Bit of Evidence,” *Economic Journal*, vol. 112(477), F107–F118.
- Podgursky, M. J. and M. G. Springer (2007): “Teacher performance pay: A review,” *Journal of Policy Analysis and Management*, vol. 26(4), 909–950.
- Rivkin, S. G., E. A. Hanushek, and J. F. Kain (2005): “Teacher, Schools and Academic Achievement,” *Econometrica*, vol. 73(2), 417–458.
- Rockoff, J. E. (2004): “The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data,” *American Economic Review*, vol. 94(2), 247–252.
- Rosen, S. (1987): “The theory of equalizing differences,” in *Handbook of Labor Economics, Handbook of Labor Economics*, vol. 1, ed. by O. Ashenfelter and R. Layard, chapter 12 (Elsevier), pp. 641–692.
- Shapiro, C. and J. E. Stiglitz (1984): “Equilibrium Unemployment as a Worker Discipline Device,” *American Economic Review*, vol. 74(3), 433–44.
- Slater, H., N. Davies, and S. Burgess (2009): “Do teachers matter?: measuring the variation in teacher effectiveness in England,” Working paper 09/212, Bristol: Centre for Market and Public Organisation.
- Stinebrickner, T. R. (1998): “An Empirical Investigation of Teacher Attrition,” *Economics of Education Review*, vol. 17(2), 127–136.

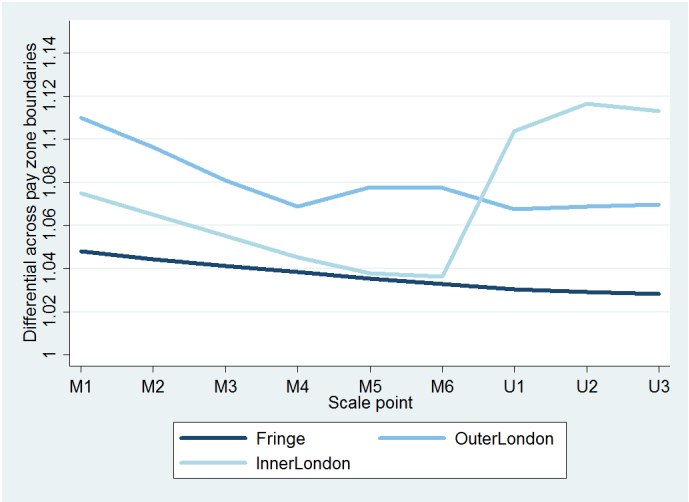
Figures and Tables

Figure 1: Teacher pay regions close to London



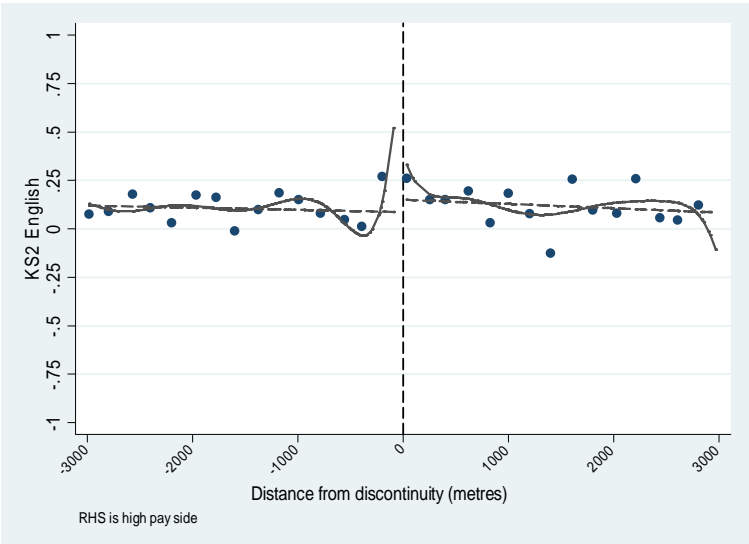
Note: Teacher pay regions around London. Primary schools are denoted by (orange) dots.
Source: School Teachers Pay and Conditions Document.

Figure 2: Differential across neighbouring pay zones



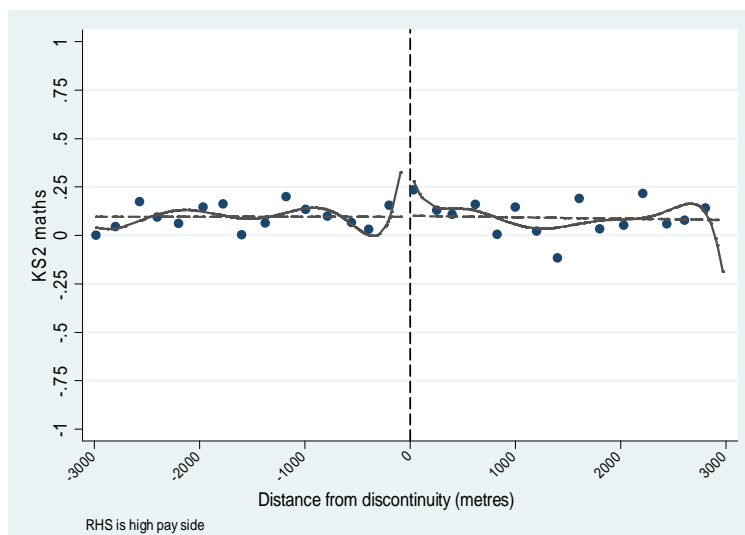
Note: Scale point refers to the teacher pay scale (ranging from lowest to highest). Source: School Teachers Pay and Conditions Document.

Figure 3: Relationship between distance to boundary and English test scores, fringe London



Note: Dashed line represents linear specification estimated either side of boundary, solid line is based on a 7th order polynomial and dots are local averages by 200m from the pay zone boundary.

Figure 4: Relationship between distance to boundary and maths test scores, fringe London



Note: Dashed line represents linear specification estimated either side of boundary, solid line is based on a 7th order polynomial and dots are local averages by 200m from the pay zone boundary.

Table 1: Summary statistics: within 2km of the boundary (all years)

Characteristic	England	London	Inner Inside	Outside	Difference	Outer Inside	Outside	Difference	Fringe Inside	Outside	Difference
Prop. FSM	0.16 (0.15)	0.21 (0.16)	0.29 (0.15)	0.20 (0.13)	0.09***	0.17 (0.13)	0.11 (0.09)	0.06***	0.08 (0.07)	0.08 (0.08)	0
Prop. SEN (no s.)	0.23 (0.11)	0.23 (0.11)	0.27 (0.11)	0.24 (0.11)	0.03***	0.23 (0.11)	0.22 (0.10)	0.02**	0.21 (0.10)	0.21 (0.11)	0
Prop. SEN (s.)	0.02 (0.02)	0.02 (0.02)	0.02 (0.02)	0.02 (0.02)	0	0.02 (0.02)	0.02 (0.02)	0.00*	0.02 (0.02)	0.02 (0.02)	0.00***
Prop. EAL	0.19 (0.24)	0.30 (0.27)	0.47 (0.23)	0.37 (0.27)	0.10***	0.20 (0.16)	0.10 (0.08)	0.10***	0.07 (0.11)	0.07 (0.10)	0
Prop. non-white	0.34 (0.30)	0.50 (0.31)	0.72 (0.22)	0.59 (0.28)	0.13***	0.39 (0.19)	0.23 (0.12)	0.16***	0.16 (0.12)	0.16 (0.13)	0
FTE Pupils (nursery to KS4)	280.6 (125.5)	310.9 (119.7)	349.7 (117.2)	351.9 (118.5)	-2.2	319.4 (121.7)	265.3 (96.1)	54.2***	259.8 (119.8)	250.2 (110.2)	9.6
IMD rank	0.54 (0.24)	0.46 (0.24)	0.27 (0.14)	0.45 (0.18)	-0.18***	0.50 (0.19)	0.62 (0.19)	-0.12***	0.73 (0.17)	0.70 (0.17)	0.03**
IMD rank (education)	0.56 (0.20)	0.56 (0.19)	0.51 (0.17)	0.58 (0.17)	-0.07***	0.53 (0.22)	0.52 (0.21)	0.01	0.62 (0.22)	0.59 (0.22)	0.03**
IMD rank (employment)	0.58 (0.22)	0.52 (0.22)	0.36 (0.15)	0.51 (0.17)	-0.15***	0.55 (0.18)	0.64 (0.19)	-0.09***	0.75 (0.16)	0.72 (0.18)	0.03**
IMD rank (income)	0.51 (0.24)	0.42 (0.24)	0.24 (0.15)	0.39 (0.18)	-0.14***	0.45 (0.20)	0.57 (0.18)	-0.12***	0.68 (0.16)	0.67 (0.18)	0.02
Prop. low IMD rank (education)	0.21 (0.27)	0.18 (0.25)	0.19 (0.25)	0.14 (0.18)	0.05***	0.28 (0.32)	0.29 (0.31)	-0.01	0.19 (0.27)	0.21 (0.26)	-0.03
Prop. low IMD rank (employment)	0.21 (0.28)	0.30 (0.30)	0.50 (0.28)	0.24 (0.22)	0.26***	0.22 (0.25)	0.13 (0.19)	0.09***	0.05 (0.13)	0.08 (0.13)	-0.03***
Prop. low IMD rank (income)	0.30 (0.34)	0.43 (0.36)	0.72 (0.27)	0.44 (0.31)	0.28***	0.36 (0.32)	0.19 (0.24)	0.18***	0.08 (0.16)	0.11 (0.18)	-0.03**
IDACI rank	0.48 (0.25)	0.38 (0.24)	0.21 (0.14)	0.35 (0.18)	-0.15***	0.40 (0.19)	0.53 (0.18)	-0.13***	0.66 (0.17)	0.65 (0.18)	0.01
KS2 total points (standardised)	0.06 (0.38)	0.05 (0.38)	-0.04 (0.36)	0.08 (0.36)	-0.12***	0.09 (0.34)	0.05 (0.36)	0.04	0.11 (0.37)	0.11 (0.36)	0
KS2 English (standardised)	0.06 (0.37)	0.06 (0.37)	-0.03 (0.35)	0.07 (0.36)	-0.10***	0.07 (0.34)	0.06 (0.35)	0.02	0.12 (0.36)	0.10 (0.36)	0.02
KS2 maths (standardised)	0.04 (0.36)	0.04 (0.37)	-0.04 (0.34)	0.07 (0.37)	-0.11***	0.08 (0.33)	0.04 (0.35)	0.04*	0.08 (0.35)	0.10 (0.34)	-0.01
Total income per pupil (log)	8.29 (0.22)	8.34 (0.22)	8.43 (0.17)	8.29 (0.17)	0.14***	8.32 (0.18)	8.19 (0.16)	0.13***	8.20 (0.18)	8.20 (0.19)	0
Total expenditure per pupil (log)	8.28 (0.22)	8.33 (0.22)	8.42 (0.17)	8.28 (0.17)	0.14***	8.31 (0.18)	8.17 (0.16)	0.13***	8.20 (0.18)	8.19 (0.19)	0
Pupil teacher ratio	21.3 (3.2)	21.6 (3.1)	21.6 (3.0)	22.1 (3.0)	-0.5***	21.3 (3.0)	22.6 (3.0)	-1.3***	21.4 (3.1)	21.3 (3.3)	0
Pupil-assistant ratio	120.6 (200.6)	111.5 (174.0)	106.0 (148.3)	119.0 (167.6)	-12.9*	105.7 (213.8)	127.8 (148.8)	-22.1*	137.1 (177.6)	119.4 (173.4)	17.7
Teacher-assistant ratio	5.5 (9.1)	5.0 (7.7)	4.8 (6.7)	5.3 (7.7)	-0.5*	4.9 (9.5)	5.5 (6.5)	-0.7	6.2 (7.9)	5.5 (8.4)	0.8
Prop. cross LA boundary	0.06 (0.10)	0.08 (0.12)	0.11 (0.13)	0.11 (0.12)	0.01*	0.09 (0.10)	0.07 (0.09)	0.01**	0.06 (0.09)	0.04 (0.09)	0.02**
Prop. cross pay region boundary	0.03 (0.07)	0.04 (0.08)	0.08 (0.11)	0.08 (0.11)	-0.01	0.06 (0.08)	0.07 (0.09)	0	0.11 (0.15)	0.08 (0.13)	0.03***
Number of obs	21912	11791	1830	1352		711	565		600	677	
Number of schools	4428	2361	366	271		143	113		120	136	

Note: * denotes significance at 10 percent; ** denotes significance at 5 percent; *** denotes significance at 1 percent. The mean characteristics for schools in England, London and within 2km of the inner London pay boundary, outer London pay boundary and fringe London pay boundary are reported. Standard deviations are reported in brackets. The sample includes all primary schools that are present in the National Pupil Database in all of the academic years 2005/2006, 2006/2007, 2007/2008, 2008/2009, 2010/2011. The sample size varies slightly between rows, where individual schools have missing data.

Table 2: Difference in Outcomes across Fringe/Rest of England Boundary in 2011: within 2km of the boundary

Outcomes	Raw difference		OLS		FILM		Matching	
	b	SE	b	SE	b	SE	b	SE
<i>Resources</i>								
Funding per pupil (log)	0.013	0.02	0.015	0.01	0.005	0.01	0.012	0.02
Pupil:Teacher Ratio	-0.027	0.38	-0.208	0.29	-0.2	0.18	-0.241	0.43
Pupil:Assistant Ratio	-5.231	6.36	-2.238	7.42	-4.458	8.97	-3.459	11.29
<i>Teacher pay</i>								
Teacher salary (£)	915.519	462.64*	475.126	274.01	369.63	293.45	720.759	590.26
Has additional responsibilities	0.032	0.02	0.013	0.02	0.016	0.02	0.018	0.03
Spine point	0.062	0.2	0.145	0.16	0.285	0.2	0.57	0.26*
<i>Pupil attainment</i>								
KS2 English (std)	0.011	0.04	0.012	0.03	0.01	0.03	-0.003	0.05
KS2 Maths (std)	-0.007	0.04	-0.014	0.03	-0.019	0.02	-0.025	0.05

Note: * denotes significance at 10 percent; ** denotes significance at 5 percent; *** denotes significance at 1 percent. The unit of observation is a primary school. The sample includes all primary schools that are present in the National Pupil Database in all of the academic years 2005/2006, 2006/2007, 2007/2008, 2008/2009, 2010/2011 and where dependent variable is observed, within 2km of a pay area boundary. The coefficient reported represents the percent increase in expenditure per pupil associated with the high-pay side of the boundary. OLS, FILM, and Matching specifications additional account for observable characteristics of the school: percentage of pupils eligible for free school meals, with English as an additional language, that are non-white, and have a special educational need, the number of pupils in the school, dummy variables for region (North-East London, South-East London, South-West London), rank of index of multiple deprivation and rank of income deprivation affecting children index. FILM and Matching additionally account for distance to the relevant boundary.

Table 3: Results: KS2 Maths Points (std): varying the measure of closeness to the Fringe boundary

Fringe/outer-fringe London boundary within 1km								
Year	Raw difference		OLS		FILM		Matching	
	b	SE	b	SE	b	SE	b	SE
2006	-0.02	0.06	-0.019	0.05	-0.032	0.03	-0.056	0.07
2007	0.053	0.07	0.036	0.07	0.019	0.03	0.019	0.07
2008	-0.017	0.06	-0.034	0.04	-0.027	0.03	-0.059	0.06
2009	0.044	0.06	0.024	0.03	0.013	0.04	-0.003	0.07
2011	-0.031	0.06	-0.003	0.05	-0.015	0.04	-0.029	0.06

Fringe/outer-fringe London boundary: within 2km								
Year	Raw difference		OLS		FILM		Matching	
	b	SE	b	SE	b	SE	b	SE
2006	-0.053	0.05	-0.049	0.02	-0.047	0.02*	-0.054	0.05
2007	-0.004	0.05	-0.017	0.03	-0.026	0.01	-0.023	0.05
2008	0.008	0.04	0.012	0.01	0.011	0.01	0.017	0.04
2009	-0.028	0.04	-0.048	0.02	-0.04	0.02*	-0.043	0.05
2011	-0.007	0.04	-0.014	0.03	-0.019	0.02	-0.025	0.05

Fringe/outer-fringe London boundary: within 3km								
Year	Raw difference		OLS		FILM		Matching	
	b	SE	b	SE	b	SE	b	SE
2006	-0.024	0.04	-0.014	0.02	-0.016	0.02	-0.016	0.04
2007	-0.003	0.04	0.001	0.02	-0.004	0.01	-0.012	0.04
2008	-0.012	0.04	0.001	0.03	-0.012	0.02	-0.008	0.05
2009	-0.045	0.04	-0.034	0.02	-0.04	0.01*	-0.041	0.04
2011	-0.039	0.04	-0.011	0.03	-0.025	0.02	-0.039	0.04

Note: * denotes significance at 10 percent; ** denotes significance at 5 percent; *** denotes significance at 1 percent. The unit of observation is a primary school. The sample includes all primary schools that are present in the National Pupil Database in all of the academic years 2005/2006, 2006/2007, 2007/2008, 2008/2009, 2010/2011 and where the dependent variable (maths scores in Key Stage 2 tests), pupil:teacher and pupil:assistant ratios are observed, within xkm of a pay area boundary. The coefficient reported represents the increase in Maths scores associated with the high-pay side of the boundary. OLS, FILM, and Matching specifications additional account for observable characteristics of the school: percentage of pupils eligible for free school meals, with English as an additional language, that are non-white, and have a special educational need, the number of pupils in the school, dummy variables for region (North-East London, South-East London, South-West London), rank of index of multiple deprivation and rank of income deprivation affecting children index. FILM and Matching additionally account for distance to the relevant boundary.

Table 4: Results: KS2 English Points (std): varying the measure of closeness to the Fringe boundary

Fringe/outer-fringe London boundary: within 1km								
Year	Raw difference		OLS		FILM		Matching	
	b	SE	b	SE	b	SE	b	SE
2006	0.018	0.07	0.006	0.05	-0.001	0.03	-0.021	0.07
2007	0.123	0.07	0.101	0.08	0.076	0.05	0.087	0.08
2008	0.004	0.06	-0.039	0.04	-0.02	0.02	-0.058	0.07
2009	0.066	0.07	0.027	0.05	0.006	0.04	-0.056	0.08
2011	-0.017	0.06	0.016	0.05	0.004	0.05	-0.021	0.07

Fringe/outer-fringe London boundary: within 2km								
Year	Raw difference		OLS		FILM		Matching	
	b	SE	b	SE	b	SE	b	SE
2006	0.002	0.05	0.01	0.03	0.018	0.01	0.01	0.05
2007	0.027	0.05	0.027	0.03	0.018	0.02	0.025	0.05
2008	0.014	0.04	0.002	0.03	0.006	0.02	0.007	0.05
2009	0.025	0.04	0.005	0.03	0	0.02	-0.022	0.05
2011	0.011	0.04	0.012	0.03	0.01	0.03	-0.003	0.05

Fringe/outer-fringe London boundary: within 3km								
Year	Raw difference		OLS		FILM		Matching	
	b	SE	b	SE	b	SE	b	SE
2006	0.01	0.04	0.023	0.03	0.025	0.01*	0.022	0.04
2007	0.006	0.04	0.009	0.02	0.005	0.02	0.003	0.04
2008	0.01	0.04	0.007	0.03	0.007	0.01	0.008	0.04
2009	-0.01	0.05	0.008	0.03	-0.013	0.02	-0.018	0.04
2011	-0.012	0.04	0.017	0.03	0.006	0.03	-0.009	0.04

Note: * denotes significance at 10 percent; ** denotes significance at 5 percent; *** denotes significance at 1 percent. The unit of observation is a primary school. The sample includes all primary schools that are present in the National Pupil Database in all of the academic years 2005/2006, 2006/2007, 2007/2008, 2008/2009, 2010/2011 and where the dependent variable (maths scores in Key Stage 2 tests), pupil:teacher and pupil:assistant ratios are observed, within xkm of a pay area boundary. The coefficient reported represents the increase in English scores associated with the high-pay side of the boundary. OLS, FILM, and Matching specifications additional account for observable characteristics of the school: percentage of pupils eligible for free school meals, with English as an additional language, that are non-white, and have a special educational need, the number of pupils in the school, dummy variables for region (North-East London, South-East London, South-West London), rank of index of multiple deprivation and rank of income deprivation affecting children index. FILM and Matching additionally account for distance to the relevant boundary.