VALUING SCHOOL QUALITY USING BOUNDARY DISCONTINUITIES

Stephen Gibbons*, Stephen Machin and Olmo Silva*****

November 2012 – Revised

* *Department of Geography and Environment and Centre for Economic Performance, London School of Economics*

** *Department of Economics, University College London and Centre for Economic Performance, London School of Economics*

*** *Department of Geography and Environment and Centre for Economic Performance, London School of Economics*

Author for correspondence:

Olmo Silva
Department of Geography and Environment,
London School of Economics,
Houghton Street, London, WC2A 2AE, UK.
Email: o.silva@lse.ac.uk
Tel.: +44 (0)20 7955 6036

VALUING SCHOOL QUALITY USING BOUNDARY DISCONTINUITIES

Abstract

Existing research shows that house prices respond to local school quality as measured by average test scores. However, higher test scores could signal higher academic value-added or higher ability, more sought-after intakes. In our research, we show that both school value-added and student prior achievement – linked to the background of children in schools – affect households' demand for education. In order to identify these effects, we improve the boundary discontinuity regression methodology by matching identical properties across admissions authority boundaries; by allowing for boundary effects and spatial trends; by re-weighting our data towards transactions that are closest to district boundaries; by eliminating boundaries that coincide with major geographical features; and by submitting our estimates to a number of novel falsification tests. Our results survive this battery of tests and show that a one-standard deviation change in either school average value-added or prior achievement raises prices by around 3%.

Keywords: House prices; school quality; boundary discontinuities.
JEL Classifications: C21; I20, H75; R21.

1. Introduction

Good schooling is frequently upheld as decisive in life, but empirical evidence remains quite ambiguous when it comes to pinning down what makes a 'good' school and what people value in education. Parents making school choices seem well aware of their preferences and go to great lengths to secure places for their children at their preferred schools. However, social scientists have had mixed success in eliciting general conclusions about the nature of these preferences.

Researchers in education have regularly used survey responses to learn about preferences for schools (e.g. Coldron and Boulton, 1991; Flatley et al., 2001; and Schneider and Buckley, 2002). The evidence from this field shows that parents rank academic outcomes highly among the reasons for choosing a school, but other factors play an important role, such as distance from home, school composition, safety and wellbeing. More recently, parents' actual choices of schools and teachers have been used as an alternative way to uncover preferences for school attributes (e.g. Hastings et al., 2005; and Jacob and Lefgren, 2007).

Apart from these examples, other research has looked for evidence of the value of schools in the capitalisation of their benefits into housing prices – i.e. using the hedonic valuation method. This wide-ranging international literature has shown that the demand for school quality is at least partly revealed in housing prices whenever school places are assigned to neighbouring homes. Gibbons and Machin (2008), Black and Machin (2010), Nguyen-Hoang and Yinger (2011) and Machin (2011) provide summaries of recent evidence, all suggesting a consensus estimate of around 3-4% house price premium for one standard deviation increase in school average test scores.

One limitation of previous work is that – with a few exceptions – it is confined to showing that prices follow headline school performance as measured by school average test scores. However, better test scores could occur through improvements in enrolment quality or through greater pupil progress – potentially driven by teaching quality, school resources, peer effects and school effectiveness. One possibility is that parents pay for school 'value-added' that represents the expected academic *gains* for their children. A second possibility is that parents pay for good peers and favourable school composition – i.e. school inputs – irrespective of the contribution of these factors to their own child's achievements.[1] While the first perspective is interesting from a policy point of view because it puts a price on interventions that raise academic standards, the second one is relevant because of its implications for school segregation (e.g. Epple and Romano, 2000).

A handful of papers have taken steps to disentangle these two channels of influence. Brasington and Haurin’s (2006) results show that that school value-added and initial achievements both have positive effects on prices, although this point is somewhat lost in their conclusions. Kane et al. (2005) also consider value-added and average test scores as alternative indicators of school performance. However, they do not present specifications that include both indicators and do not aim to provide evidence on the importance of value-added. In contrast, Clapp et al. (2007) show that pupil ethnicity seems more important than test scores to home buyers around Connecticut schools, although the authors do not have access to data on pupils’ academic progress.

Other papers have looked at the importance of school expenditure relative to test score outputs. For example, Downes and Zabel (2002) find that test scores are capitalised into local house prices, whereas measures of school expenditures are not. Cellini et al. (2010) use referenda outcomes in California’s school finance system to suggest that house prices respond to the level of capital expenditure per pupil and that this cannot be fully explained by changes in test scores. Occasionally other school attributes have been considered. For example, Figlio and Lucas (2004) find that state-assigned school ratings have a transient effect on prices, over and above test scores, suggesting that householders draw additional information about achievement from these grades, or else value the ratings in their own right. Finally, Gibbons and Machin (2006) suggest that popularity in itself

[1] Kramarz et al. (2009) provide empirical tests of the relative importance of pupil, school and peer effects in determining test scores. Their findings suggest that a large part of the variation in test scores is explained by pupil attributes, followed by school quality differentials, while peers’ characteristics matter less. This is consistent with Gibbons and Telhaj (2008), Lavy et al. (2012) and most other studies of peer effects.

raises prices, given that over-capacity schools command an additional premium relative to under-capacity schools with equal performance.

Our paper moves this literature forward in a number of important ways. Our first contribution is to use a convincing strategy to show that house prices respond causally to school age-7 to age-11 test score gains (value-added), indicating that parents value school educational output. Our results suggest that parents also value the average age-7 test score component of this value-added measure, which we interpret as a marker for students' background characteristics. We argue that this result arises from parental demand for good school composition, rather than demand for school quality in the early years, even if school composition is not a productive input in the educational production function. This interpretation is supported by further evidence showing that the price effects from age-7 achievements are completely explained by students' background characteristics, especially their eligibility for free meals (a proxy for low family income).

Our second contribution is to further refine, improve and test the boundary discontinuity regression method, which is the 'state-of-the-art' approach used in this field to mitigate potential biases induced by neighbourhood unobservables. We present several innovations and refinements, which can be summarised as follows: *(a)* We combine matching methods with the regression-discontinuity design to allow for flexibility in the way in which housing observables affect price differentials across boundaries; *(b)* We incorporate in our models a variety of boundary fixed effects and spatial trends to account semi-parametrically for between-district unobserved heterogeneity and trends in amenities across boundaries; *(c)* We inverse-distance weight our regressions such that identification comes from variation at the admission zone boundaries where neighbourhood heterogeneity is minimised; this refines previous studies which used samples restricted to fixed buffer-zones close to boundaries (e.g. 1/4 mile); *(e)* We perform a number of falsification exercises and a compelling placebo test which uses the quality of autonomous state schools that do *not* admit on the basis of residential location, but administer the same standard tests as the mainstream schools that prioritise admission on place of residence.

A final advantage of our work is that we establish these findings using large scale administrative data for the whole of England, and not just for one city (e.g. Boston or San Francisco) as done by much of the previous research. The size and coverage of our data makes the above strategies feasible and the findings more representative.

To preview our results, we find that a one-standard deviation change in either age-7 to age-11 school average value-added or prior (age 7) achievement raises prices by around 3% for schools that prioritise students who live close by. Conversely, we show that there is no house price premium attached to properties close to high quality schools that do *not* prioritise local students. This finding – alongside other falsification exercises – demonstrates that our findings are causal and not spurious.[2] Finally, various back-of-the envelope calculations show that the magnitude of this house price response to school quality is plausible as a parental investment decision given the expected return in terms of future earnings of their children.

The remainder of the paper has the following structure. Section 2 explains our methods. Section 3 discusses the context in which we apply our approach and the data setup. Section 4 presents our results and discussion. Finally, Section 4.7 concludes.

2. Empirical Strategy

2.1. Methodological framework

Our empirical work uses a geographical boundary-based regression discontinuity design. This approach was initially popularised by the work of Black (1999), with several more recent examples (e.g. Bogart and Cromwell, 2000; Gibbons and Machin, 2003, 2006; Bayer and McMillan, 2005; Kane et al., 2005; Davidoff and Leigh, 2007; Fack and Grenet, 2010; Bayer et al., 2007, and Ries and Somerville, 2010). Closely related studies investigate the effects of local taxes (Cushing, 1984,

---

[2] Note that this is different from the exercise of Fack and Grenet (2010), who show that house prices respond less to the quality of local non-autonomous schools if there are autonomous schools in the area. The authors cannot perform a similar falsification test because their autonomous schools are private schools and are not ranked using comparable performance tables as state schools (unlike ours) .

Duranton et al., 2006; Holmes, 1998) and market access when there are changes in national borders and their permeability (Redding and Sturm, 2008; Hanson, 2003).

The standard hedonic property value model (Sheppard, 1999) represents property market prices (usually log prices) as a linear combination of observable property attributes and the implicit market price of these attributes. The implicit prices can be estimated by standard least squares regression techniques, but researchers usually do not observe all salient property and neighbourhood characteristics, leading to omitted variable biases. This problem is particularly acute for amenities – e.g. school performance – that depend on the distribution of characteristics in the local population, and hence on sorting in relation to unobserved area effects.

A way to mitigate this problem is to difference the data between close-neighbouring houses to eliminate area-specific unobservables, but this strategy only works for school quality if there is a sharp discontinuity in its supply between close-neighbouring homes. This condition holds when admissions involve contiguous pre-defined admission zones such that residents on each side of the boundary have access to different sets of schools. Regression specifications can then include attendance district boundary dummy variables, or be estimated on data that is differenced between matched pairs of neighbouring houses on either side of the boundary. This research design is set out below in a way that will help explain our methods.

The price p (in logs) of a house sale, with characteristics $x(c)$ in a location c, is:

$$p = s(c)\beta + x(c)\gamma + g(c) + \varepsilon \tag{1}$$

where $s(c)$ represents the expected school 'quality' that residents can access at residence at c. 'Quality' includes both school composition and effectiveness, and in our empirical application we estimate the effects of these different components separately. As usual, ε represents unobserved housing attributes and errors that are assumed to be independent of x and c. The function $g(c)$ represents unobserved influences on market prices that are correlated across neighbouring locations, such that the price varies deterministically with geographical position. Location c can be specified in various ways, most flexibly in terms of a vector of geographical or Cartesian coordinates.

2.2. Identification issues in geographical boundary discontinuity models

The fundamental identification problem arises because of the common dependence of prices, housing characteristics and anticipated school quality on $g(c)$, which can be potentially eliminated by spatial differencing between locations i and j:

$$(p_i - p_j) = (s(c_i) - s(c_j))\beta + (x_i(c_i) - x_j(c_j))\gamma + g(c_i) - g(c_j) + (\varepsilon_i - \varepsilon_j) \quad (2)$$

and choosing i and j to be as geographically close as is feasibly possible. Consistent estimation of the implicit prices (β, γ) requires the unobservable spatial component $g(c_i) - g(c_j)$ to be effectively random, and (conditional on *x*) uncorrelated with the difference in school quality $s(c_i) - s(c_j)$. This condition will not hold in general, and will require the researcher to find locations i, j such that locally $Cov\left[s(c_i) - s(c_j), g(c_i) - g(c_j)\right] = 0$ and $Var\left[s(c_i) - s(c_j)\right] \neq 0$ (conditional on observed housing and neighbourhood characteristics). These two conditions are unlikely to be met for any continuous functions $s(.), g(.)$ because the first requires that $c_i = c_j$, which would violate the second. However, the two conditions can hold approximately for closely spaced neighbours if $s(.)$ is discontinuous and $g(.)$ is continuous such that:

A1: $Var\left[g(c_i) - g(c_j)\right] \to 0$ as $|c_i - c_j| \to 0$, where $|c_i - c_j|$ is the Euclidian distance between house sales i and j.

A2: $Var\left[s(c_i) - s(c_j)\right] \to \theta$ as $|c_i - c_j| \to 0$, where θ is a positive constant (or positive definite matrix if *s* is multidimensional).

The geographical boundary discontinuity approach exploits A1 by choosing i, j to be as close together as possible, whilst ensuring that i, j are on different sides of an attendance zone boundary to satisfy A2. Note that the geographical boundary discontinuity method differs from the standard regression discontinuity design (Imbens and Lemieux, 2008) in which a single forcing variable (e.g. voting share, as in Lee et al., 2004) determines treatment (e.g. party affiliation of an elected representative), although the general principle is similar.

In practical empirical settings, there are three main threats to the identification strategy sketched above:

(a) There are spatial trends in amenities across boundaries, implying that even if assumption A1 holds in principle, it is violated in practice because the distance between sales $|c_i - c_j|$ in housing sales samples is never exactly zero. Such cases occur when spatially correlated amenities cause house prices on one side of a boundary to differ *on average* from those on the other side (highly localised factors contained in $\varepsilon_i - \varepsilon_j$ are not a concern). Examples of these include a rail station or a good *secondary* school driving up prices in one direction moving away from a *primary* school admissions boundary (we will return to this issue in Section 3.2), pulling richer families and better school intakes into the high price side.

(b) There are boundary discontinuities in prices not caused by school quality differences, but driven by a discontinuity in $g(c)$, violating assumption A1. A typical case occurs when attendance zone boundaries coincide with geographical features (roads, railways etc.) that partition communities. Other cases can arise without visible evidence of the boundary on the ground, if houses on different sides have different directional aspect or outlook (e.g. one south facing, the other north facing), or because districts have different tax rates or district-specific local services, like refuse collection or policing.

(c) School quality is not discontinuous at the boundary, violating assumption A2. This could occur if attendance boundaries are porous allowing pupils to attend schools in neighbouring districts because of historical changes to school admissions policies or attendance zones not being enforced. Note however, that the discontinuity can be ‘fuzzy’ and identification requires only that there is a change in the probability of attending different sets of schools, and hence expected school quality, as one crosses the boundary.

2.3. Methods to address the identification problems

Although the concerns discussed above have been partly addressed in the literature, our analysis goes much further in establishing the credibility of the boundary discontinuity approach through a series of robustness and falsification tests. Some of these checks are refinements of methods that have been previously applied, while others are completely novel. For clarity, these strategies are labelled M1-M7 and this coding is used in the results section.

M1. *Difference between matched property transactions with identical observable characteristics across administrative boundaries*. Following the literature on non-parametric discrete-cell matching (Rubin, 1973), we pair up each house sale with the nearest transaction of the same property type on the opposite side of an administrative attendance district in the same year (Fack and Grenet, 2010, use a similar idea but match to fictitious predicted sales prices).

M2. *Weight regressions to zero-distance housing transaction pairs and include distance to boundary polynomials to control for cross-boundary price trends of the type (a) discussed above.* Earlier work (e.g. Black, 1999) tested robustness to cross-boundary trends by restricting the analysis to narrow distance buffers along the boundary, i.e. applying weights of 1 to transactions within the buffer and zero otherwise. We generalise this by weighting observations in inverse proportion to the distance between sales, such that greater weight applies to observations that are close neighbours (on opposite sides of the boundary). This is an important contribution of our approach given that conditions A1 and A2 hold as the distance between paired transactions approaches zero. Following the regression discontinuity literature, we also control for polynomial trends in distance from the boundary discontinuity (e.g. DiNardo and Lee, 2004; Lee et al., 2004; and Clark, 2009).

M3. *Include boundary fixed effects in cross-boundary difference models to control for boundary-specific discontinuities of the type (b) discussed above*. Our institutional context features multiple schools on each side of an attendance district boundary, so that school quality varies across boundaries *and* along a boundary within a given attendance district. This data structure allows for boundary fixed effects in our cross-boundary differenced model, thus eliminating between-boundary variation due to unobservable factors fixed along to the boundary. This is crucial given assumption A1 and the problems with boundary-specific discontinuities highlighted in Section 2.2 under case (b).[3]

M4. *Restrict our attention to boundaries where pupils rarely cross to address fuzzy discontinuities of the type (c) discussed above*. Our data allows us to observe whether pupils cross an admission district boundary to attend their school, so we can check the sensitivity of our results to the fuzziness of the school quality discontinuity arising from pupil flows across boundaries.

M5. *Eliminate boundaries that coincide with significant geographical obstacles*. Our analysis uses only inland school district boundaries that do not coincide with tidal estuaries and rivers (e.g. the Thames in London). In addition, we eliminate portions of the boundaries that coincide with major roads, motorways and railways, to test sensitivity to non-schooling related sources of price discontinuity as in point (b) above.

M6. *Compare results for cases in which home location is and is not a school admission criterion*. Our institutional context has two types of schools. For the majority (65%) non-autonomous institutions, school places are typically allocated first to pupils who live closest to the school and attendance district boundaries are binding. There are therefore compelling reasons to buy a home close to a school of choice and on the 'good' side of the boundary. On the other hand, autonomous schools (mainly religious, comprising about 35% of schools) operate pupil admissions policies that *do not* compel families to buy their home close to the school (e.g.

[3] This approach is different from Dhar and Ross (2012) who exploit time variation to include boundary effects.

based on church attendance and denomination). Although parents might still buy a house close to the school of choice to minimise travel costs, they do *not need* to do so to secure admission to their children. Thus, we expect local house prices to respond to the quality of non-autonomous schools, but not to the quality of autonomous schools. This institutional feature provides us with a novel and demanding falsification test based on the comparison of the price response to the quality of both types of schools as an additional check on the issues raised in points (a) and (b) in Section 2.2. We discuss these features of the school admission system in more detail in Section 3.2.

M7. *Apply falsification tests using fake attendance boundaries*. We re-estimate our models using differences between transactions in our boundary sample that are separated by similar distances to the transactions matched across boundaries, but actually lie *within* the same attendance district. This method was first applied in Black (1999) for Boston, MA. However, we go further and provide a powerful falsification test using differences between property transactions along a network of imaginary attendance boundaries that mimics real patterns, but is geographically translated in a south-westerly direction. A finding of a positive association between school quality and housing prices in this case would falsify the claim that price effects are causally linked to cross-boundary school quality discontinuities. This exercise addresses the concerns raised in point (a) in Section 2.2 and tests assumption A1.

The robustness and falsification tests described above relate to identification of the causal effect of school quality and other characteristics on house prices. However, there is an additional set of identification issues that arise when households' preferences and/or incomes are heterogeneous and the aim is to interpret regression estimates as measures of the 'willingness to pay' for school quality. With sorting of households in response to school quality, a hedonic regression does not necessarily estimate the *mean* valuation of school quality, because willingness to pay varies with household characteristics. Bayer et al. (2007) (building on methods introduced by Berry et al., 1995) provide a structural solution to this problem, but their method relies on strong and hard-to-test assumptions. We do not wish to impose this much structure and present no novel solution to this problem. Therefore, in the presence of heterogeneous preferences and sorting across boundaries, our discontinuity design will provide a weighted average of the marginal WTP of residents along the admissions zone boundary. In our defence, the work by Bayer et al. (2007) shows that traditional hedonic regressions *are* effective at evaluating mean WTP in contexts (like ours) where the amenity in question is supplied at various qualities in many different locations.[4]

A second consequence of sorting is that it becomes difficult to separate willingness to pay for school quality from willingness to pay for neighbours' quality. Part of the association between school quality and house prices necessarily works through its effect on neighbours' quality. Our robustness checks in this respect are limited to a control variable strategy in which many of the neighbourhood demographic controls are potentially endogenous. Nevertheless, our evidence shows that our estimates of the value of school quality are steadfastly linked directly to school attributes, not to neighbourhood quality. In addition, we graphically and statistically test for the presence of discontinuities in salient area characteristics following the regression discontinuity design literature (and similar to Bayer et al., 2007, and Kane et al., 2005), and find none.

3. Institutional Context, Data and Empirical Specification

3.1. National curriculum and assessment in England

Compulsory education in England is organised into five stages referred to as Key Stages. In the primary phase, pupils enter school at age 4-5 in the Foundation Stage then move on to Key Stage 1 (*ks1*), spanning ages 5-6 and 6-7. At age 7-8 pupils move to Key Stage 2, sometimes – but not usually – with a change of school.[5] At the end of Key Stage 2 (*ks2*), when they are 10-11, children leave the primary phase and go on to secondary school where they progress through Key Stage 3

[4] The authors find a house price response of approximately 2.5% for a one standard deviation change in test scores in their 'standard' hedonic models, which rises to around 3% when accounting for the effects of sorting.

[5] In some cases there are separate Infants and Junior schools (covering Key Stage 1 and 2 respectively) and a few LAs operate a Middle School system (bridging the primary and secondary phases) – we do not consider these schools here.

and 4. At the end of each Key Stage, in May, pupils are assessed on the basis of standard national tests, and progress through the phases is measured in terms of Key Stage Levels, ranging between W (working towards Level 1) and Level 5+ in the primary phase. A point system can also be applied to convert these levels into scores that represent about one term's (10-12 weeks) progress.

Since 1996, in the autumn of each year, the results of the National Curriculum assessment at Key Stage 2 are published as a guide to primary school performance. More recently, since 2003, a value-added score has also been reported, based on the average pupil gain at each school between age 7 and age 11 (relative to the national average). Schools and Local Education Authorities report these performance figures in their admissions documents, and parents refer to these documents and the performance tables, as well as using word-of-mouth recommendations, when choosing schools (see, *inter alia*, Flatley et al., 2001 and Gibbons and Silva, 2011b).

In our empirical work below, we use the *ks1* to *ks2* test score value-added (*va*) as the main indicator of schools' production *output*, or effectiveness, while *ks1* scores serve as a control for pupils' prior academic achievements. These *ks1* scores provide a summary indicator of school *inputs* embodied in the pupil intake, which parents can infer from the published *ks2* scores, school visits, word of mouth, and other local knowledge.[6] These *ks1* tests could reflect the effectiveness of a school in a child's early years, but they are not made public and so cannot provide parents with a direct signal of school performance. In Section 4.6 we directly test and discuss the significance of other pupil background characteristics as proxies for school intake quality.

3.2. School types, admissions and boundaries

All state primary schools in England are funded largely by the central government, through Local Authorities (LAs, formerly Local Education Authorities) that are responsible for schools in their

[6] Performance tables contain information on the fraction of students with special education needs (SEN), with varying degrees of severity. SEN status is partly based on poor performance in early tests and assessments. Thus parents can gather some indirect information about the intake quality of a school using performance tables. More recently the tables provide much more contextual data on student background, but this information postdates the periods in our study.

geographical domain. These schools fall into a number of different categories and differ in terms of their governance and who controls pupil admissions.[7] Two thirds of primary schools are termed 'Community' schools and are closely controlled by the LA. The remaining third are usually linked to a Faith or charitable organisation and are more autonomously run. The key difference relevant to this paper is between schools that make their own choices on whom to admit – which we term *autonomous* schools – and *non-autonomous* schools such as Community schools to which pupils are assigned by the Local Authority. Gibbons et al. (2008) and Gibbons and Silva (2011a) provide more details on the overall differences between these two groups of schools.

All LAs and schools must organise their admissions arrangements in accordance with the current (now statutory) School Admissions Code. The guiding principle is that parental choice should be the first consideration when ranking applications to a primary school. However, if the number of applicants exceeds the number of available places, almost any criterion, which is not discriminatory, does not involve selection by ability and can be clearly assessed by parents, can be used to prioritise applicants. These criteria vary in detail and change over time, but preference in non-autonomous schools is usually given first to children with special educational needs, next to children with siblings in the school and, crucially, to those children who live closest. For Faith and other autonomous schools, regular attendance at church and other expressions of religious commitment are foremost priorities. Place of residence almost never features among admission guidelines for Faith schools, and when it does it relates to Diocese boundaries, which do not follow non-autonomous school admission boundaries. Consequently, there is little reason for parents to pay for homes close to good autonomous schools, other than to reduce travel costs.

Another crucial feature of the admission system that applies to non-autonomous, but not to autonomous primary schools, is that pupils hardly ever attended non-autonomous primary schools

[7] LAs are responsible for the strategic management of state education services, including planning the supply of school places, intervening where a school is failing and allocating central funding to schools. Additionally, there is a small private, fee-paying sector, educating around 6-7% of pupils in England, which we do not consider here.

outside of their LA of residence during the years under analysis. This is because, although families are allowed to apply to non-autonomous schools in other LAs, up until recently parents had to make separate applications to different LAs. More importantly, LAs did not have a statutory requirement to find a school for pupils from other school districts.[8] As a result, banking on admission to a popular non-autonomous primary school in another LA is a high-risk strategy, and LA boundaries effectively act as primary school admissions district boundaries over the period we study. This provides a source of discontinuity in the non-autonomous primary school 'quality' that residents can access on different sides of LA boundaries.

Note also that LA barriers are not binding for secondary school admissions (age 12 onwards), which are fewer in number (approximately 2500, compared to around 14000 primary schools), attract students over greater distance (nearly three times the median distance for primary schools) and, crucially, from different LAs.

3.3. Source data

Three main data sources are used. The Price-paid dataset from the UK Land Registry for the years 2000-2006 provides administrative data on address, sales price and basic characteristics of all domestic properties sold in the UK. Each property is located by its postcode – typically 17 neighbouring addresses – and each postcode is assigned to a 1-metre geographical coordinate.

Data on each pupil's assessment record for all Key Stages plus information on school attended, gender, age, ethnicity, language skills, special educational needs, entitlement to free meals, residential postcode are obtained from the National Pupil Database (NPD). This is an administrative pupil-level dataset collected by the UK's Department for Education (DfE). Additional information on schools is gathered from the "Edubase" database, the Annual School Census and the Consistent

---

[8] More precisely, the Education Act 1996 section 14 reads: "(1) A Local Education Authority shall secure that sufficient schools for providing (a) Primary education, and (b) education that is Secondary education (…) for their area. (2) The schools available for an area shall not be regarded as sufficient (…) unless they are sufficient in number, character and equipment to provide for all pupils the opportunity of appropriate education".

Financial Reporting series that can also be obtained from the DfE. Neighbourhood characteristics at Output Area (OA) level come from the 2001 GB Census, and are linked to the housing transactions data by postcode.[9]

3.4. Linking schools to housing transactions and matching across boundaries

Linking the school data to housing sales is complex, because the admissions system implies that there is no one-to-one mapping between a postcode and the set of schools accessible from that location *within* the LA of residence. We infer this mapping using a computationally intensive, but intuitive approach. This procedure imputes the set of schools accessible from each postcode in our Price-paid housing transactions database using the actual attendance patterns of pupils that are recorded in the NPD. This is much more sophisticated than just assigning a house to a set of nearest schools, and is essential when we want to exploit discontinuities generated by LA boundaries. Moreover, this 'revealed preference' definition of catchment areas implicitly accounts for features of school choice and attendance patterns that would be obscured by simpler assignment rules. The exact mapping procedure, its justification and the practical implementation are described in detail in Appendix A.

After creating each school-specific catchment area, we calculate the distance and direction from each school to each housing transaction postcode (up to a maximum distance of 10km – the 99.5th percentile of home-school distance at age 7) and assign each postcode to multiple schools by deducing which housing transactions lie within which school catchment areas. Postcode-specific school characteristics are then obtained by averaging the characteristics of the schools to which the postcode is assigned and applying higher weights to the nearest schools (unweighted means give similar results).

The procedure described above yields a dataset of over 1.6 million transactions for 2003, 2004, 2005 and 2006 joined to data on the average characteristics of the schools that are accessible from

[9] OAs in England contain on average 130 households and less than ten postcodes.

the postcode of each sale. To set up the spatially-differenced cross-boundary model in Equation (2) property transactions are matched to the nearest transaction in the same year of the same property type (detached, semi-detached, terraced or flats, and leasehold or freehold) occurring in an adjacent LA (method M1).[10] For method M5, we further cut out boundaries that coincide with major roads, motorways and railways, while for method M7 we set up a set of matched sales across fake LA boundaries and a set of matched sales within LAs. For the within-LA matching, we pair houses in the boundary sub-sample imposing the constraint that the matched sale is within the same LA and at least 20m away (to achieve better comparability with the cross-LA samples).To generate the fake boundary matching, we simply translate the geographical coordinates of the housing transactions data by 10km North and 10km East, and repeat the matching exercise.

3.5. Empirical specification

Applying the data described above to Equation (2) yields the following empirical specification:

$$\Delta p_{hi} = \beta_1 \Delta va_i + \beta_2 \Delta ks1_i + \Delta z_i' \lambda + \Delta x_{hi}' \gamma + \Delta g\left(c_i\right) + \Delta \varepsilon_{hi} \tag{3}$$

where p_{hi} is the (log) price of the house sale h in postcode i. The notation Δ means a difference between matched, closest transactions on either side of an LA boundary. The variable va_i is the expected value-added and $ks1_i$ is the mean age-7 test score (our marker for background and prior achievement) for schools that can be accessed from location i. The vector z_i contains other observable school and neighbourhood characteristics; the vector x_h contains observable attributes of house sale h; and the function $\Delta g\left(c_i\right)$ is parameterised using boundary dummy variables, distance to school, distance between matched transactions and various distance-to-boundary polynomials as set out in M2 and M3. Finally, ε_i represents the usual error term.

---

[10] Matching is restricted initially to properties with 2.5km, and then to within the median inter-property distance along the boundary so that the maximum distance varies with the density of transactions.

Although we have house sales and school attributes in multiple periods, we have suppressed the *t*-subscripts in Equation (3) for simplicity. The pupil census occurs in January, while pupils take their *ks1* and *ks2* assessments in May and the results are published towards the end of the calendar year. We therefore link prices of houses sold in calendar year *t* (January to December) to the age-7 to age-11 valued-added and school census figures published at the end of year *t-1* (October to November). Hence, *ks1* results relate to the period *t-5*, although we present additional results that include *ks1* measured just prior to the transaction date at time (*t-1*). Note that we do not exploit the time dimension alone in our identification strategy by differencing over time because test scores assigned to house postcodes are highly serially correlated and short-run changes are likely to be uninformative about school quality (Kane and Staiger 2002). Thus, in the next section we present results from regression estimates of the models in Equation (3) obtained by pooling all periods.

4. Results

4.1. Descriptive statistics

Table 1 presents some key descriptive statistics. The first two columns summarise the full data set, while the second two summarise the boundary sub-sample lying within 2500m of an LA border (described above). The average price of sales in the full data is £182,730. In the boundary sub-sample the mean is about £13,000 or 7% higher. This is because administrative boundaries are more prevalent in and around towns and cities and there is a greater chance of finding matched pairs of sales in densely populated areas along boundaries in urban areas. Figure 1 illustrates the general spread of sales throughout England following the administrative boundary structure by showing the locations of transactions in the boundary sub-sample for two arbitrarily chosen geographical areas: the Midlands, North West and South Yorkshire (*Panel A*); and London and the South East (*Panel B*).

In terms of school test scores, value-added is higher in the boundary sub-sample and *ks1*scores lower, but the differences are small. Houses in this sub-sample have slightly fewer accessible schools and are closer to schools. This is consistent with our claim that LA boundaries restrict the choice set for houses located close to the boundary (see the discussion above and in Gibbons et al., 2008), and with the urban nature of the sub-sample. For the boundary group, we also present statistics on the distance to the closest boundary and the distance between property pairs that are matched across boundaries. The raw mean distance to the boundary and nearest matched property is around 500 metres and 735 metres respectively. In the regressions, we apply inverse inter-sale distance weights, so the inverse distance weighted (IDW) means provide a better representation of the effective distances. These weighted distances are only 133m to the boundary and 206m to the nearest matched property.

4.2. Evaluating the boundary discontinuities

To support our argument that the LA district boundary creates a barrier to school admissions, we show that cross-district school attendance is much less prevalent than within-district attendance close to district boundaries. The relevant figures are presented in Table 2 and refer to proportions of students in the postcode. In the full dataset, only 3.3% of pupils attend schools other than in their home LAs, which largely reflects the fact that schools in neighbouring LAs will be further away. In the boundary sub-sample the proportion rises to 6.2%, while the IDW mean proportion crossing from each residential postcode in our sales data is 25%. Since this figure corresponds to addresses only 133m from the boundary (Table 1), we would expect a nearly 50% chance of attending a school on either side of the boundary if this did not impose a barrier and was unimportant for admission. Moreover, these means are from distributions that are highly right-skewed: the median proportion of pupils attending a school in a different district is zero.

Explicit tests for discontinuities in school quality and other area characteristics at the LA boundary are provided in Figure 2 and Figure 3. In all these plots, the *x*-axis reports the distance to the LA boundary. The right hand side of the diagram (distance > 0) corresponds to sales which have access to greater school value-added than their match across the boundary, i.e. $s(c_i)-s(c_j)>0$ in

Equation (2). The left side of the diagram (distance < 0) corresponds to postcodes with access to value-added below that on the other side of the boundary. The figures are obtained as predictions from a regression of the cross-boundary difference in the relevant variable on a positive side and negative side constant term, and distance-decile dummies (up to 800m) from the boundary on each side. The dependent variables are standardised by the standard deviation of the cross-boundary difference (within 800m). The dotted lines show 95% confidence intervals based on standard errors clustered on the matched postcode c_j. The plots are restricted to 500m on each side for clarity and shown alongside a test for whether the differences on both sides at the boundary are equal (i.e. an F-test of the hypothesis that the coefficients on the 'good' and 'bad' side dummies are equal).

The top left panel of Figure 2 shows a large discontinuity in value-added scores at LA boundaries for non-autonomous schools (Assumption A2). This large difference occurs by construction given the way the right and left halves of the plot are defined. More importantly, almost half of the 2-standard deviation spread occurs within the first 100m, from where our identification will predominantly come. The top right panel shows that a discontinuity in house sale prices exists too: although visually small, the difference across the boundary is highly significant. Comparing the top left and right panels suggests that a 0.8 standard deviation change in school average value-added is associated with a 0.05 standard deviation change in house prices at the boundary. Note that prices do not follow school average value-added as the distance to the boundary increases because other amenities drive these spatial price trends. This illustrates the importance of weighting our regression estimates to close-neighbour observations and controlling for distance-to-boundary trends (method M2).

In Figure 3 we present similar pictures for a range of neighbourhood-related characteristics, with left and right sides split by low and high non-autonomous school value-added. These plots show whether cross-boundary neighbourhood differences are correlated with cross-boundary non-autonomous school value-added differences. Evidently there are no large discontinuities in any of the salient demographic or socioeconomic characteristics we consider. In addition, there are no significant cross-boundary differences in the average distance to schools or the number of schools accessible from a postcode, allaying concern about travel costs or the degree of choice differing on opposite side of LA borders.

4.3. Baseline results: comparing the price effects of school value-added and prior achievements

Table 3 presents coefficients and standard errors for our main regression results. We report only the key figures for school-mean value-added ('output') and *ks1* test scores ('composition'). The reported coefficients are multiplied by 100 to show, to an approximation, the percentage effect of a one point change in school mean test scores. Control variables are listed in the table notes. The specifications become increasingly stringent as we move left to right across the table. Column (1) reports results from a simple OLS regression using the full time-pooled cross-sectional samples for 2002-2006; Column (2) shows the same specification estimated on the boundary sub-sample and Column (3) presents the cross-boundary (method M1) pair-wise differenced model described in Section 2.3 and in Equation (3). Columns (4) to (7) introduce the other modifications described in Section 2.3, by adding inverse distance weighting and distance-to-boundary polynomials (M2), LA boundary dummies (M3), by focussing on boundaries with below-median rates of crossing (M4) and by eliminating boundaries that coincide with geographical features (M5).

To begin with, we discuss the price effects of value-added. Value-added is obtained as the difference between age-11 students' *ks2* data published just prior to the sale, and the *ks1* scores from these same students when aged 7 four years earlier. In the OLS estimates, we observe very large and significant associations between school value-added and house prices with a one point change linked to an 11-14% change in prices (8-11% for a one standard deviation change in the school average value-added distribution). Evidently these estimates are not causal: as soon as we apply the boundary differencing strategy there is a dramatic fall in the price effect of school value-added, down to 2% in Column (3). After applying inverse distance (IDW) weights in Column (4) and adding distance-to-boundary polynomials, the coefficient on value-added rises to 3.8% and becomes more statistically significant (alternative distance-penalising weighting schemes gave similar results). Note that if we follow Black (1999) and only concentrate on the closest properties pairs (e.g. we apply equal weights of 1 to all transactions 250m apart and zero otherwise) we find similar results.

An important result is that once we apply IDW weights, the coefficient on value-added remains very stable even when we add in boundary dummy variables as in Column (5). We can further include boundary × year dummies to eliminate all time-series variation occurring along boundaries and the value-added coefficient is almost unchanged (at 3.74). All in all, our specifications indicate that prices rise by about 3.7-3.8% for a one point increase in school value-added from the mean, or about 3% for a one standard deviation change in the school average value-added distribution.

Our results also show a positive relationship between age-7 test (*ks1*) scores and house prices. These *ks1*scores are those used to construct student value-added, i.e. they are tests taken four years earlier. The OLS results on the full sample show a 3.7% change in prices for a one point change in *ks1* test scores. Once we focus our attention to the boundary sample and apply IDW weights, the effect is reduced, but remains significant and suggests a price response of around 2.8% for a one point improvement or about 3% for a one standard deviation change. The interpretation we place on this coefficient is that it measures the house price response due to parental demand for peer quality, irrespective of its impact on test score progression. Comparing the response to value-added and age-7 scores, it is evident that school choice is driven by the demand both for expected academic gain *and* for aspects of peer group quality that are uncorrelated with current academic gains (i.e. school intake composition conditional on school value-added). The net result is that house prices respond to mean age-11 (*ks2*) test scores, whether or not these arise through school composition or school value-added.

An alternative interpretation is that *ks1* achievements measure pre-age 7 school value-added, although a number of factors count against this interpretation. Firstly, cross-boundary differences in pupil background characteristics – i.e. free-meal entitlement, ethnicity and special educational needs – account for 40% of the variance in cross-boundary differences in *ks1* achievements. On the other hand, cross-boundary differences in value-added and cross-boundary differences in *ks1* only share 2.9% of their variance (i.e. the square of the correlation between *va* and *ks1* is 0.029). Therefore, pupil background is the main observable component in the variation of the *ks1* test scores included in the regressions reported in Table 3. Secondly, as we will show in Section 4.6, the effect of *ks1* is eliminated when we include school intake characteristics in our specification. In short, age-7 test scores are most likely proxying for the background of the school intake, and not for pre-age 7 teaching quality and school effectiveness in general. We will return to this point in Section 4.6 and in our subsequent Discussion.

Our results change only slightly when we test their sensitivity to the fuzziness in the discontinuity by restricting our analysis to boundaries with a sharp discontinuity due to low rates of crossing (below median, or less than 5% of pupils crossing along the whole boundary) as in Column (6). Similarly, eliminating cases where the boundaries might split communities because they coincide with major roads, motorways or railways makes very little difference (Column (7)). The results are also robust to other specifications (not reported in the table) which include: interactions between distance-to-boundary and boundary dummies; two year averages of test scores to reduce downward biases from idiosyncratic noise in the single year variables (Kane et al., 2003 and Gibbons and Machin, 2003); and splitting of the sample into groups with high and low *ks2* dispersion (across schools) to check whether uncertainty over where a child will go to school matters.

4.4. Falsification tests using schools that do not admit pupils based on home location

Method M6 discussed above proposed a novel falsification test based on testing whether house prices respond to the quality of schools that do *not* ration places according to home address. Some initial visual evidence is provided in Figure 4, which is analogous to Figure 2, but depicts the cross-boundary discontinuity for autonomous school quality, where the right hand side corresponds to places with relatively high *autonomous* school quality (and vice versa for the left hand side). By construction there is a strong rise in school quality across the boundary (left hand panel). However, there is no discontinuity in house prices at the boundary (right hand panel) suggesting that there is no price effect from the quality of schools that do not prioritise admissions to pupils living close by.

Table 4 presents some corresponding regression results where we compare the effect of school quality on house prices for autonomous and non-autonomous schools. The first two rows present the association of house prices with quality in non-autonomous schools that admit pupils according to home address. The second two rows show the coefficients for autonomous schools for which home-to-school distance and LA of residence are not important admission criteria. In the OLS estimates presented in Columns (1), we find that the association between school quality and housing prices is large and significant for both types of school. Given that the only reason to buy a property close to an autonomous school is to minimise transport costs – not to grant admission – these associations between autonomous school quality and house prices most likely reflects a reverse-causal relationship between local family incomes (driven by differences in neighbourhood amenities, such as access to better transport) and average academic achievement in schools that pupils from these families attend. In contrast, as soon as we difference across LA boundaries as in Column (2), we still find positive and significant results for non-autonomous schools, but very small and insignificant effects for autonomous schools. A joint test for the coefficients on value-added and age-7 test scores in Column (2) being equal for autonomous and non-autonomous schools clearly rejects the null hypothesis with a p-value of 0.025.

A further concern is that shrewd parents seeking admission to popular autonomous schools might buy cheaper housing on the sides of the LA boundaries with low non-autonomous school quality and then send their children across the boundary to an autonomous school. This could attenuate our estimates of the effect of non-autonomous school, with autonomous school quality lifting housing prices when non-autonomous quality is low. However, as shown in Column (3), the interactions between autonomous and non-autonomous school quality are not significantly linked to prices either, making this hypothesis highly unlikely.

4.5. Falsification tests using fake and inoperative boundaries

In Table 5, we implement the falsification tests presented as method M7 and based on 'imaginary' boundaries. In the first two columns we pair sales in our boundary sub-sample with sales *within* the same LA. A similar test was carried out in Black (1999). Column (1) presents the simple OLS estimates, while Column (2) presents the coefficients based on the differenced data. OLS estimates are similar to what we found before on the boundary sub-sample. However, when we difference between close-neighbour pairs within the same LA, we find no house price effects. This suggests that our main findings are not spuriously driven by local unobservables, but causally linked to cross-boundary school quality discontinuities. Pairing transactions across fake boundaries translated in a south-westerly direction as in Columns (3) and (4) yields similar results. In summary, these two tests do not falsify our claim that the findings in Table 3 are causal estimates arising from the demand for school quality when admission is constrained by *real* attendance boundaries.

4.6. Robustness of the results to sorting, neighbourhood attributes and other school inputs

Sorting of heterogeneous households across boundaries could lead to biases – especially a spurious association between house prices and *ks1* scores – if prices respond to differences in neighbourhood quality rather than school quality, and sorting leads to differences in school composition. To account for these possible issues, we introduce a variety of (potentially endogenous) neighbourhood demographic and school-level control variables in our regressions in Table 6.

Column (1) repeats our preferred specification from Table 3. Column (2) adds in *ks1* scores of age-7 pupils in the school just before (at time *t-1*) the time of the house purchase (at time *t*), alongside the *ks1* scores of the pupils on whom our value-added measure (at *t-1*) was based. We find that this control for recent *ks1* achievements is insignificant and makes no difference to our main coefficients. This result suggests that: *(i)* parents pay attention to the characteristics of pupils who have just completed primary school and for whom *ks2* test results are available, rather than to the background of younger pupils in the school; and *(ii)* the estimated effect of *ks1* scores is causal, and does not arise as a result of reverse causation running from house price to recent school intake differentials. The next two columns add in a set of Census demographic control variables described in the table notes (Column (3)), or the average school achievements of children in the neighbourhood (this is defined as a geographical area that shares the same three nearest schools; Column (4)). In either case, there is little change in the coefficients on school quality, particularly school value-added, which confirms that school effectiveness is capitalised into house prices over and above the characteristics and educational progress of pupils living in the same neighbourhood. Overall, these findings suggest that residential sorting has little bearing on our estimates of the price effects of school performance – especially the contribution of value-added.

School financial resources also have a potential relationship with housing prices – through taxes and through family background linkages – and this is an issue that we have not discussed yet. In England, resources are allocated to LAs from central government grant on the basis of needs (mainly numbers of pupils, levels of income, disadvantage and special educational needs). However, LAs tend to distribute this grant to their schools simply on the basis of pupil numbers, with various other small payments and allowances for severe special educational needs (Sibieta et al., 2008). Most of the variation in school expenditure per pupil is therefore between-LAs and hence taken out by our LA-pair boundary dummies (method M3). It is however possible that localised factors within LAs (e.g. parents' fund raising associations) generate correlation between within-LA expenditure per pupil and within-LA house prices. Therefore, we add to our specifications controls for school resources (pupil teacher ratio, expenditure per pupil and pupil numbers) along with a control for local housing tax rates (Column (5)). These variables show no statistically significant association with prices, and the main results on school quality remain unchanged.[11] Next, in Column (6), we include school demographic characteristics that affect school income, namely percentages of pupils eligible for free meals, from ethnic minorities and with special educational needs. In this specification, the coefficient on *ks1* test scores falls to near zero and is statistically insignificant. This is mainly because the school proportion of low-income pupils eligible for free meals does a better job at measuring those dimensions of school composition that influence parental demand and house prices. Other aspects of school composition – ethnicity and special educational needs – turn out to be irrelevant. In contrast, although the coefficient on value-added is attenuated slightly in this saturated model, it remains highly statistically significant and important in size, emphasising the crucial role of value-added in driving the house price response.

4.7. Interpretation and discussion

Our results show that households pay significant house price premium to gain access to schools that are likely to raise their children's educational achievements – i.e. high value-added schools. The results also suggest that households pay an additional premium for a favourable distribution of pupil characteristics – which we represented by higher mean achievements at age 7. This premium seems to be linked to the willingness of households to pay for a more favourable family income distribution in the school – namely, fewer children on free school meals – rather than school effectiveness at the earliest stages of education. Other factors such as ethnic mix, higher school expenditures or smaller classes do not influence demand (conditional on value-added).

As it turns out, the magnitudes of the effects of school composition and value-added are similar, which implies that a one point increase in school average test scores at age 11 is valued approximately the same, irrespective of whether this is achieved through value-added or school composition. One potential explanation is that parents use the headline, end-of-primary test scores

[11] This is not surprising given the weak link between resources and performance that can be observed within cross-sectional data on state school systems (see among others Hanushek, 2003, and Levacic and Vignoles, 2002).

as an indicator of academic effectiveness, but do not differentiate between test scores generated by school effectiveness and those due to a school enrolling high achieving pupils from the start. An implication of this conjecture is that households are paying *in part* for aspects of schools that are unlikely to make much difference to their own child's achievement. Another possibility is that value-added is really just another dimension of school composition, reflecting the average rate of progress of pupils enrolling in a school, but unrelated to the expected gains the school would generate for a child picked at random. The implication then is that parents pay to access schools that admit fast-progressing pupils, even though these schools offer no obvious academic benefits to their own child. Both these scenarios seem theoretically and empirically unappealing (see Section 4.6). The most plausible explanation that is consistent with our results is that parents value both academic effectiveness and composition aspects of school quality because they are interested in their own child's academic progress, as well as the social status of their child's peers.

5. Conclusion

A principal objective of this paper was to establish whether the well-documented response of housing prices to school-mean test scores represents a demand for the educational value-added output of schools, or demand for components of school quality that are desirable, but unlikely to raise a child's achievements. Therefore, our first research aim was to go further than previous work in finding out if, why, and by how much people pay for homes near good schools. This is a crucial policy question, because if prices respond to educational output, this signals a value in public investments in teaching quality, leadership and resources, that potentially lead to higher achievements, better lifetime outcomes and economic performance.

In order to tackle this issue, we developed a number of refinements to the boundary discontinuity approach and applied a series of novel robustness checks and falsification tests in order to establish causality. These methods are of broader interest in that they potentially generalise to other contexts such as border effects in international trade (Redding and Sturm, 2008; Hanson, 2004), provision of health care (Cooper et al., 2011), and the effects of local tax regimes and policies on housing costs and business location (Cushing, 1984; Holmes, 1998).

Our headline finding is that house prices respond equally to both the expected academic gains during primary school and the initial characteristics of students. Most importantly, the statistical association between school value-added and house prices appears to be causal: our estimates are empirically indestructible, regardless of our many attempts to falsify this claim by testing for alternative causal channels. This finding persuades us that parents really do care about value-added when they value schools.

The magnitude of our estimates of the effect of school quality is in line with previous research for England and internationally (see Gibbons and Machin, 2008, Black and Machin 2010): prices increase from the mean by about 3% for a one standard deviation improvement in school-mean age-7 to age-11 value-added, plus about 3% for a one standard deviation increase in mean school achievements at age 7. It is interesting to note that these figures are plausible when benchmarked against alternative options and the future labour market returns. A standard deviation in the pupil test score distribution is worth around 11% or £20,500 on house prices, equivalent to just over 2.5 years of private schooling fees at the time of this study. Further, the present value of the labour market returns of this educational improvement – at approximately £20,600 – are of a similar order of magnitude to its capitalised housing value.[12]

[12] See Appendix B for details of these calculations.

Acknowledgements: We would like to thank Amy Challen and Anushri Bansal for excellent research assistance, and the Editor, two anonymous referees, Jaap Abbring, Victor Lavy, Erik Sørensen, and participants at numerous seminars and conferences for comments and suggestions. We are responsible for any errors or omissions.

References

Bayer, P., F. Ferreira and R. McMillan (2007): "A Unified Framework for Measuring preferences for Schools and Neighborhoods", *Journal of Political Economy*, vol. 115, 588-638.

Bayer, P. and R. McMillan (2005): "Choice and Competition in Local Education Markets", NBER Working Paper 11802.

Berry, S., J. Levinsohn and A. Pakes (1995): "Automobile Prices in Market Equilibrium", *Econometrica*, vol. 63, 841-890.

Black, S. (1999): "Do Better Schools Matter? Parental Valuation of Elementary Education", *Quarterly Journal of Economics*, vol. 114, 577-599.

Black, S. and S. Machin (2010): "Housing Valuations of School Performance", E. Hanushek, S. Machin and L. Woessmann (eds.) *Handbook of the Economics of Education*, Amsterdam: Elsevier.

Bogart, W. and B. Cromwell (2000): "How much is a Neighbourhood School Worth?", *Journal of Urban Economics*, vol. 47, 280-305.

Brasington, D. and D. Haurin (2006): "Educational Outcomes and House Values: A Test of the Value Added Approach", *Journal of Regional Science*, vol. 46, 245-268.

Cellini, S., F. Ferreira and J. Rothstein (2010): "The Value of School Facilities: Evidence from a Dynamic Regression Discontinuity Design", *Quarterly Journal of Economics*, vol. 125, 215-61.

Clapp, J., A. Nanda and S. L Ross (2008): "Which School Attributes Matter? The Influence of School District Performance and Demographic Composition on Property Values", *Journal of Urban Economics*, vol. 63, 451-466.

Clark, D. (2009): "The Performance and Competitive Effects of School Autonomy", *Journal of Political Economy*, vol. 117, 745-783.

Coldron, J. and P. Boulton (1991): "Happiness' as a Criterion of Parents Choice of School", *Journal of Education Policy*, vol. 6, 169-178.

Cooper, Z., S. Gibbons, S. Jones, A. McGuire (2011):"Does hospital competition save lives? Evidence from recent English NHS reforms", *Economic Journal*, vol. 121(554), F228–F260.

Cushing, B. (1984): "Capitalization of Interjurisdictional Fiscal Differentials: An Alternative Approach", *Journal of Urban Economics*, vol. 15, 327-356.

Dhar, P. and S. Ross (2012): "School district quality and property Values: Examining differences along school district boundaries", *Journal of Urban Economics*, vol. 71, 18-25.

Davidoff, I. and A. Leigh (2006): "How Much Do Public Schools Really Cost? Estimating the Relationship

Between House Prices and School Quality", Australian National University, mimeo.

DiNardo, J. and D. Lee (2004): "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001", *Quarterly Journal of Economics*, vol. 119, 1383-1441.

Downes, T. and J. Zabel (2002): "The Impact of School Characteristics on House Prices: Chicago 1987-1991", *Journal of Urban Economics*, vol. 52, 1-25.

Duranton, G., L. Gobillon, H. Overman (2006): "Assessing the Effects of Local Taxation Using Micro Geographic Data", London School of Economics, CEP Discussion Paper 748.

Epple, D. and R. Romano (2000): "Neighborhood schools, choice, and the distribution of educational benefits", NBER Working Paper 7850.

Fack, G. and J. Grenet (2010): "When Do Better Schools Raise Housing Prices? Evidence from Public and Private Schools", *Journal of Public Economics*, vol. 94, 59-77.

Figlio, D. and M. Lucas (2004): "What's in a Grade? School Report Cards and the Housing Market", *American Economic Review*, vol. 94, pp. 591-604.

Flatley, J. H., V. Connolly, J. Higgins, J. Williams, K. Coldron, K.A. Stephenson, A. Logie and N. Smith (2001): "Parents' Experiences of the Process of Choosing a Secondary School", DCSF, London, Research Report 278 at http://www.dfes.gov.uk/research/data/uploadfiles/RR278.PDF

Gibbons, S. (2003): "Paying for Good Neighbours: Estimating the Value of an Implied Educated Community", *Urban Studies*, vol. 40, 809-833

Gibbons, S. and S. Machin (2003): "Valuing English Primary Schools", *Journal of Urban Economics*, vol. 53, 197-219.

Gibbons, S. and S. Machin (2005): "Valuing Rail Access Using Transport Innovations", *Journal of Urban Economics*, vol. 57, 148-169.

Gibbons, S. and S. Machin (2006): "Paying for Primary Schools: Admissions Constraints, School Popularity or Congestion", *Economic Journal*, vol. 116, C77-C92.

Gibbons, S and S. Machin (2008): "Valuing School Quality, Better Transport and Lower Crime: Evidence from House Prices", *Oxford Review of Economic Policy*, vol. 24, 99-119.

Gibbons, S, S. Machin and O. Silva (2008): "Competition, Choice and Pupil Achievement", *Journal of the European Economic Association*, vol. 6, 912-947.

Gibbons, S. and O. Silva (2011a): "Faith Primary Schools: Better Schools or Better Pupils?", *Journal of Labor Economics*, vol. 29, 589-635.

Gibbons, S. and O. Silva (2011b): "School Quality, Child Wellbeing and Parents' Satisfaction", *Economics of Education Review*, vol. 30, 312-331.

Gibbons, S. and S. Telhaj (2008): "Peers and achievement in England's secondary schools", SERC Discussion Paper 1.

Hanson, G. (2004): "What Has Happened to Wages in Mexico since NAFTA? Implications for Hemispheric Free Trade", in A. Estevadeordal, D. Rodrik, A. Taylor and A. Velasco (Eds.), *Integrating the Americas: FTAA and Beyond*, Harvard University Press, 505-538.

Hanushek, E. (2003): "The Failure of Input-Based Policies", *Economic Journal*, vol. 113(485), F64-98.

Hanushek, E., J. Kain, J. Markman and S. Rivkin (2003): "Does Peer Ability Affect Student Achievement?", *Journal of Applied Econometrics*, vol. 18, pp. 527-544.

Hastings, J., T. Kane and D. Staiger (2005): "Parental Preferences and School Competition: Evidence from a Public School Choice Program", NBER Working Paper 11805.

Hobbs, G. and A. Vignoles (2010): "Is Free School Meal Status a Valid Proxy for Socio-Economic Status in Schools Research?", British Educational Research Journal, 36, 673-90.

Holmes, T. (1998): "The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders", *Journal of Political Economy*, vol. 106(4), 667-705.

Hoxby, C. (2002): "Peer Effects in the Classroom: Learning from Gender and Race Variation", NBER Working Paper 7867.

Imbens, G. and T. Lemieux (2008): "Regression Discontinuity Designs: A Guide to Practice", *Journal of Econometrics*, vol. 142, 615-635.

Jacob, B. and L. Lefgren (2007): "What Do Parents Value in Education? An Empirical Investigation of Parents' Revealed Preferences for Teachers", *Quarterly Journal of Economics*, vol. 122, 1603-1637.

Kane, T. and D. Staiger (2002) "Volatility in School Test Scores: Implications for Test-Based Accountability Systems", In D. Ravith (Ed.), *Brookings papers on education policy*, 2002, pp. 235–260. Washington, DC: Brookings Institution

Kane, T, D. Staiger and S. Riegg (2005): "School Quality, Neighbourhoods and Housing Prices: The Impacts of Desegregation", NBER Working Paper 11347.

Kane, T, D. Staiger and G. Samms (2003): "School Accountability Ratings and Housing Values", *Brooking Papers on Urban Affairs*, 83-137.

Kramarz, F., S. Machin and A. Ouazad (2009): "What Makes a Test Score? The Respective Contributions of Pupils, Schools and Peers in Achievement in English Primary Education", CEE Discussion Paper 102.

Lavy, V., O. Silva and F. Weinhardt (2012): "The Good, the Bad and the Average: Evidence on Ability Peer Effects in Schools", *Journal of Labor Economics*, vol. 30, 367-414.

Lee, D., E. Moretti and M. Butler (2004): “Do Voter Affect or Elect Policies? Evidence from the U.S. House”, *Quarterly Journal of Economics*, vol. 199, 807-859.

Levacic, R. and A. Vignoles (2002): “Researching the Links between School Resources and Student Outcomes in the UK: A Review of Issues and Evidence”, *Education Economics*, vol. 10(3), 313-331.

Machin, S. (2011): "Houses and Schools - Valuation of School Quality Through the Housing Market", *Labour Economics*, 18, 723-29.

Machin, S. and S. McNally (2008): “The Literacy Hour”, *Journal of Public Economics*, vol. 92, 1441-1462.

Machin, S., S. McNally and C. Meghir (2010): “Resources and Standards in Urban Schools”, *Journal of Human Capital*, 4, 365-93.

Nguyen-Hoang, P. and J. Yinger (2011): “The Capitalization of School Quality into House Values: A Review”, *Journal of Housing Economics*, vol. 20, 30-48.

Redding, S. and D. Sturm (2008): “The Costs of Remoteness: Evidence from German Division and Reunification”, *American Economic Review*, vol. 98, 1766-97.

Ries, J. and T. Somerville (2010): “School Quality and Residential Property Values: Evidence from Vancouver Rezoning”, *Review of Economics and Statistics*, vol. 92, 928-944.

Rubin, D. (1973): “Matching to Remove Bias in Observational Studies”, *Biometrics*, vol. 29, pp. 159-183.

Sibieta, L., H. Chowdry and A. Muriel (2008): “Level Playing Field? The Implications of School Funding”, IFS Press Release and CfBT Education Trust Research Paper.

Schneider, M. and J. Buckley (2002): “What Do Parents Want from Schools? Evidence from the Internet”, *Education Evaluation and Policy Analysis*, vol. 24, 133-144.

Sheppard, S. (1999): “Hedonic Analysis of Housing Markets”, in P. Cheshire and E. Mills (Eds.), *Handbook of Regional and Urban Economics, Volume 3: Applied Urban Economics*, Amsterdam, North Holland, 1595-1635.

6. Tables

Table 1: Descriptive statistics

	Full data set		Boundary sub-sample	
	Mean	s.d.	Mean	s.d.
Price	182730	153372	195910	165360
Log price	11.91	0.642	11.98	0.625
Age 11-7 value-added	12.60	0.789	12.69	0.781
Age 7 English and Maths points	14.90	1.093	14.62	1.087
Age 11 English and Maths points	27.50	1.235	27.31	1.189
Number of schools in catchment area	3.98	2.19	3.871	1.937
Distance from home to school	2289.4	1376	1779.5	1083.8
Distance to boundary	-	-	492.6	347.4
Inverse distance weighted distance to boundary	-	-	133.2	202.9
Distance between properties	-	-	723.1	402.2
Inverse distance weighted property distance	-	-	205.5	133.2
Observations	1656056		138132	

Table 2: Statistics for pupils crossing admission district boundaries

	Full data set	Boundary sub-sample
Mean postcode proportion non-autonomous boundary crossers	0.033	0.062
IDW mean postcode proportion non-autonomous crossers	-	0.250
Median postcode proportion non-autonomous boundary crossers	-	0

Notes: Figures refer to proportions in the postcode. IDW means weighted by inverse distance between matched property transactions pairs (i.e. weighted toward observations that have zero-distance matches on opposite side to admission district boundary).

Table 3: OLS and cross-boundary difference models of the effect of school quality measures on house prices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	OLS all England	OLS boundary	Cross-LA boundary	Cross-LA boundary	Cross-LA boundary	Low cross. sample	Eliminate geo-features
Method:			M1	M2	M3	M4	M5
Age 11-7 Value-added	**10.64	**14.23	**2.06	**3.81	**3.69	**3.49	**3.62
(year t-1 – t-5)	(0.55)	(1.03)	(0.52)	(0.90)	(0.87)	(1.09)	(0.95)
Age7 English, maths	**3.66	0.53	**3.57	**2.87	**2.75	**3.07	*2.27
(year t-5)	(0.45)	(1.05)	(0.52)	(0.85)	(0.80)	(0.91)	(0.91)
Inverse distance weights & boundary distance cubics	No	No	No	Yes	Yes	Yes	Yes
Admissions authority boundary fixed effects	No	No	No	No	Yes	Yes	Yes
Observations	1656001	138132	138132	138132	138132	60394	118779

Notes: Table reports regression coefficients and standard errors multiplied by 100 to give the % effect of a one point change in explanatory variables. Dependent variable: log house sales price. School characteristics imputed from schools accessible from housing transaction site. Control variables are: average rooms per dwelling in transaction's census 2001 output area, census output area proportion of households social renting, census ward population density, ward proportion under continuous or semi-continuous urban land cover, number of schools accessible from transaction site, average distance to accessible schools, distance from transaction site to local authority boundary, year dummies. Sample based on transaction pairs for second-hand home sales in years 2003, 2004, 2005 and first quarter of 2006, from Land Registry "Price-paid" postcode dataset. Columns (1) and (2) include additional controls for property type (detached, semi-detached, terraced, flat/maisonette) and ownership type (leasehold or freehold). All variables in Columns (3) to (7) are differenced between neighbouring transaction pairs on opposite sides of school admissions authority boundary, where neighbouring pairs are matched by transaction year, property type and ownership type. Column (6) sample restricted to boundaries with below-median proportions (<5%) of pupils crossing. Column (7) eliminates cases where boundaries coincide with major roads, motorways and railways. Standard errors are clustered on matched nearest sites across boundaries (15489 clusters, Columns (3) to (7)), or clustered on Census ward (Columns (1) and (2)). Test for equality of coefficients on age 7 tests and value-added in weighted x-LA models Column (4) to (7) fails to reject null (e.g.: Column (6), p-value = 0.359).

Table 4: Falsification checks with autonomous schools (Method M8)

	(1)	(2)	(3)
	OLS on boundary sample	Cross-LA boundary M6	Cross-LA boundary M6
Age 11-7 Value-added (year t-1 – t-5), non-autonomous schools	**14.46 (1.03)	**3.68 (0.87)	**3.70 (0.87)
Age7 English, maths (year t-5), non-autonomous schools	-1.23 (1.16)	**2.72 (0.80)	**2.72 (0.80)
Age 11-7 Value-added (year t-1 – t-5), in autonomous schools	**9.89 (1.05)	0.72 (0.80)	0.74 (0.89)
Age7 English, maths (year t-5), autonomous schools	**5.76 (0.97	0.70 (0.80)	0.66 (0.80)
Age 11-7 value-added autonomous x autonomous	-	-	1.93 (1.15)
Age 7 English maths, autonomous x autonomous	-	-	-0.63 (0.83)
Inverse distance weights and distance to boundary cubic	No	Yes	Yes
Admissions boundary dummies	No	Yes	Yes
Observations	138132	138132	138132

Notes: As for Table 3 and 4. Column (1) includes additional controls for property type (detached, semi-detached, terraced, flat/maisonette) and ownership type (leasehold or freehold). All variables in Columns (2) to (3) are differences between neighbouring transaction pairs on opposite sides of school admissions authority boundary, where neighbouring pairs are matched by transaction year, property type and ownership type. Standard errors are clustered on matched nearest sites across boundaries (15489 clusters, Columns (2) to (3)), or clustered on Census ward (Columns (1)).

Table 5: Falsification tests: Within-admissions zone and fake boundary difference models of the effect of school quality on house prices (Method M7)

	(1)	(2)	(3)	(4)
	OLS within-LA boundary sample	Within-LA boundary sample	OLS fake boundary sample	Cross fake LA boundary
Age 11-7 Value-added	**14.96	0.55	**16.85	0.57
(year t-1 – t-5)	(0.94)	(0.54)	(1.50)	(1.56)
Age7 English, maths	**3.28	0.79	-0.328	0.15
(year t-5)	(0.83)	(0.48)	(1.83)	(1.23)
IDW & boundary distance cubic	No	Yes	No	Yes
Admissions boundary dummies	-	-	No	Yes
Observations	130500	130500	92054	92054

Notes: As for Table 3. Column (1) includes additional controls for property type (detached, semi-detached, terraced, flat/maisonette) and ownership type (leasehold or freehold). All variables in Columns (2) and are differences between neighbouring transaction pairs on same side of school admissions authority boundaries, where neighbouring pairs are matched by transaction year, property type and ownership type, and a minimum distance of 20m and maximum distance of 1500m is imposed. Variables in Column (4) are differences between neighbouring transaction pairs on opposite sides of 'fake' school admissions authority boundaries, where neighbouring pairs are matched by transaction year, property type and ownership type. Fake boundaries are created by translation 10km North and East. Standard errors are clustered on matched nearest sites (Columns (2) and (4)), or clustered on Census ward (Columns (1) and (3)).

Table 6: Some models with additional (potentially endogenous) controls

	(1) Cross-LA boundary	(2) Cross-LA boundary	(3) Cross-LA boundary	(4) Cross-LA boundary	(5) Cross-LA boundary	(6) Cross-LA boundary
Age 11-7 Value-added (year t-1 – t-5)	**3.69 (0.87)	**3.52 (0.89)	**3.11 (0.88)	**3.86 (1.01)	**3.18 (0.85)	**2.37 (0.87)
Age7 English, maths (year t-5)	**2.75 (0.80)	*2.25 (0.93)	*1.82 (0.79)	**2.47 (0.80)	**2.38 (0.76)	0.37 (0.85)
Age7 English, maths (year t-1)	No	0.94 (0.79)	No	No	No	No
Neighbourhood control variables	No	No	p=0.000	p=0.000	p=0.000	p=0.000
House neighbourhood Age 7-11 value-added and age 7 scores	No	No	No	p=0.006	No	No
School expenditure	No	No	No	No	p=0.385	No
Local housing (council) tax rate	No	No	No	No	Yes	No
Pupil characteristics	No	No	No	No	No	p=0.069
Standard controls	Yes	Yes	Yes	Yes	Yes	Yes
IDW & boundary distance cubic	Yes	Yes	Yes	Yes	Yes	Yes
Admissions authority boundary effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	138132	130972	138120	104929	137118	137643

Notes: As for Table 3. All variables are differences between neighbouring transaction pairs on opposite sides of school admissions authority boundary, where neighbouring pairs are matched by transaction year, property type and ownership type. Standard errors are clustered on matched nearest sites across boundaries. Neighbourhood control variables include proportions high qualified, unqualified, black, labour market active, unemployed, with dependant children, retired and of homes sold. School expenditure and local taxes control set includes expenditure per pupil, pupil-teacher ratio, number of full-time equivalent pupils and local housing taxes. Pupil characteristics include percentage of pupil eligible for free school meals, percentage of pupils from ethnic minority and percentage of pupils with special educational needs. Neighbourhood characteristics are measured at Census Output Area level, the smallest geographical unit in the GB 2001 Census containing on average 130 households.

Figure 1: Example extracts from the boundary sample

Panel A: Map of the Midlands, Manchester and Yorkshire

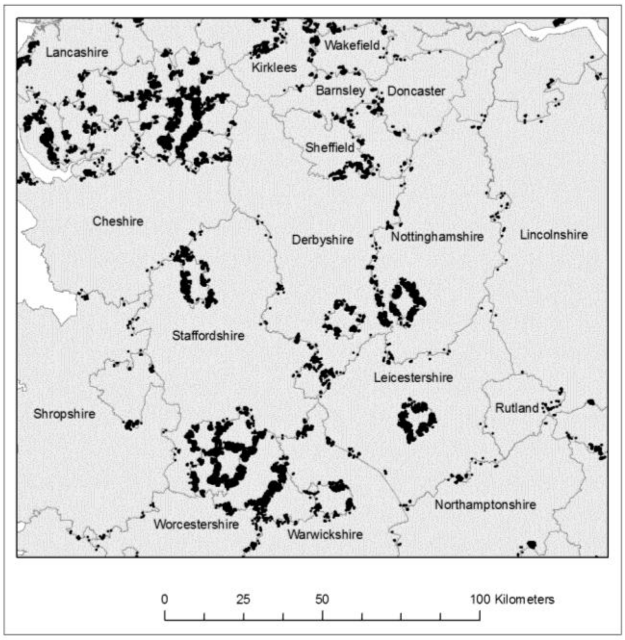

Panel B: Map of the London and the South-East

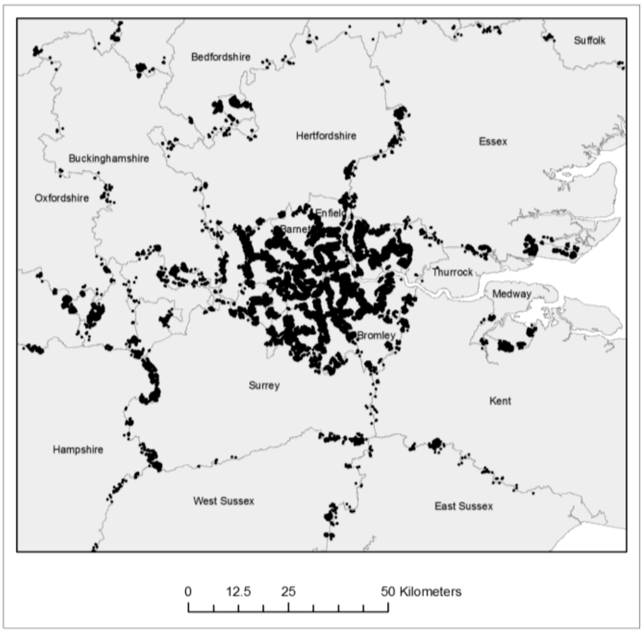

Note: The data used in the empirical analysis covers boundaries over all of England.

Figure 2: Discontinuities in non-autonomous school quality and house prices

Non-autonomous value-added, by non-autonomous value-added, p=0.000

Log house price, by non-autonomous value-added, p=0.006

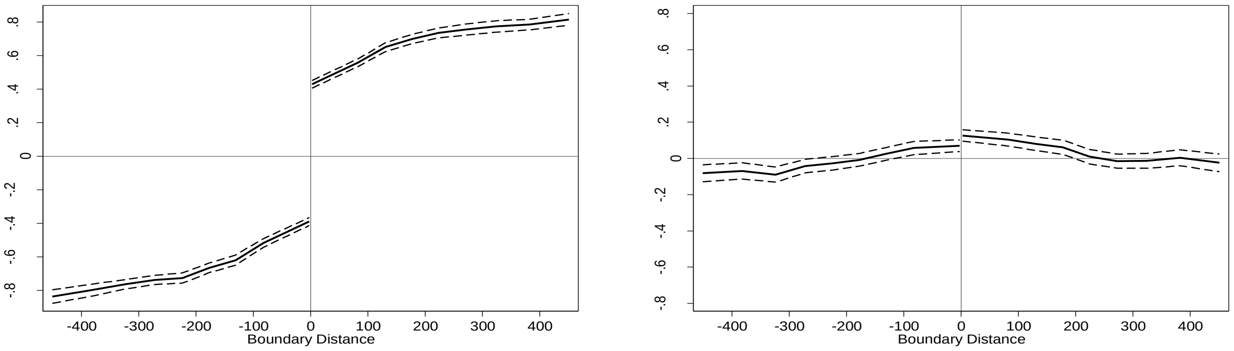

Notes: The scale on the x-axis is in metres from the boundary, at the minimum of each bin used in the regressions. The scale on the y-axis is in standard deviations.

Figure 3: Discontinuities and non-discontinuities in neighbourhood characteristics

Dwelling size. p=0.429

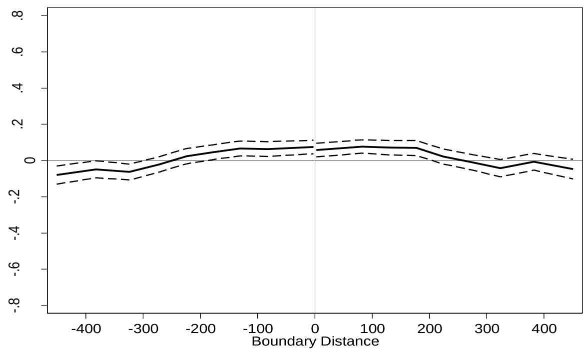

Population density: p=0.147

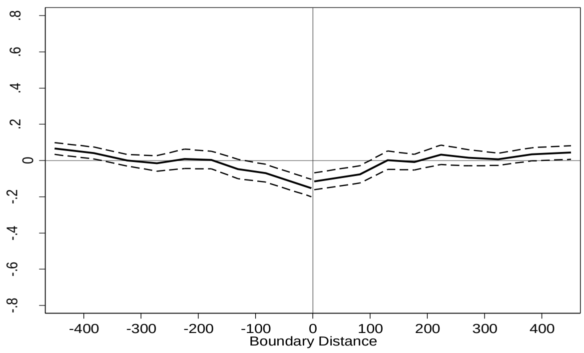

Proportion black, p=0.152

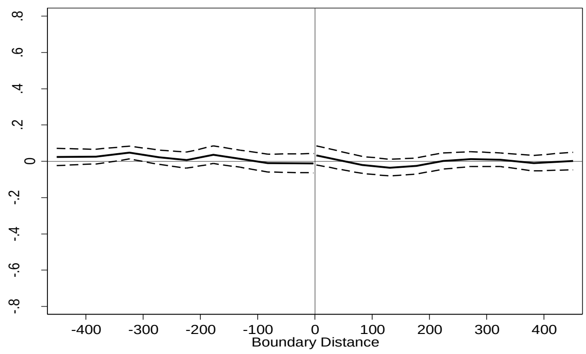

Share of dwellings sold, p=0.076

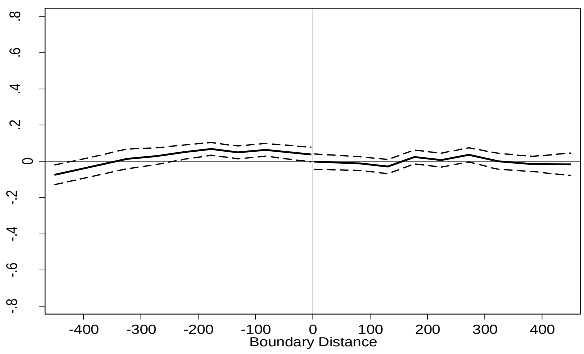

Proportion social tenants, p=0.100

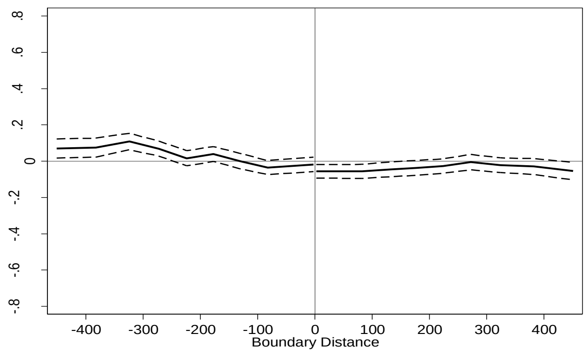

Proportion unemployed, p=0.598

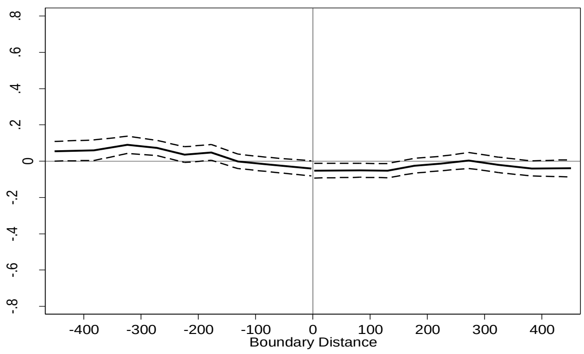

Figure 3 (Continued)

Proportion labour market active, p=0.266

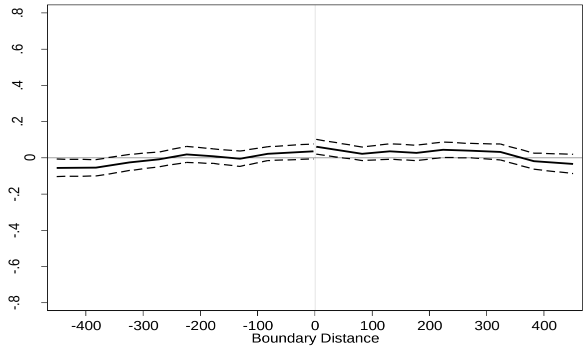

Proportion high qualified, p=0.118

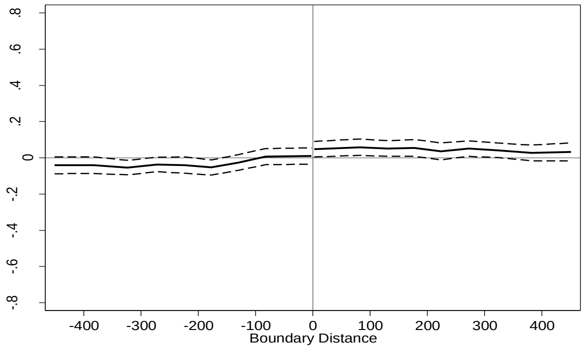

Average distance to schools in catchment area, p=0.902

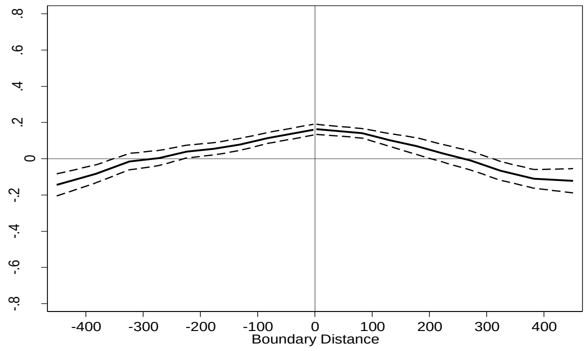

Number of schools in catchment area, p=0.572

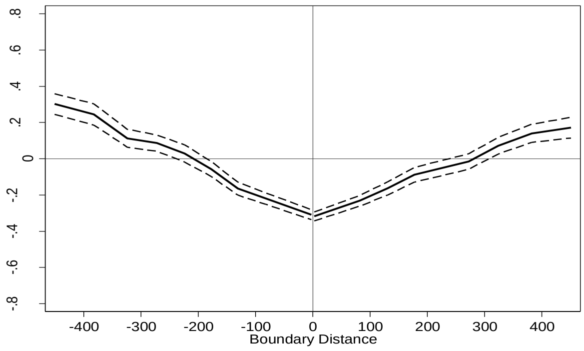

Notes: The scale on the x-axis is in metres from the boundary, at the minimum of each bin used in the regressions. The scale on the y-axis is in standard deviations.

Figure 4: Discontinuities and non-discontinuities in autonomous school quality and house prices

Autonomous value-added, by autonomous value-added, p=0.000

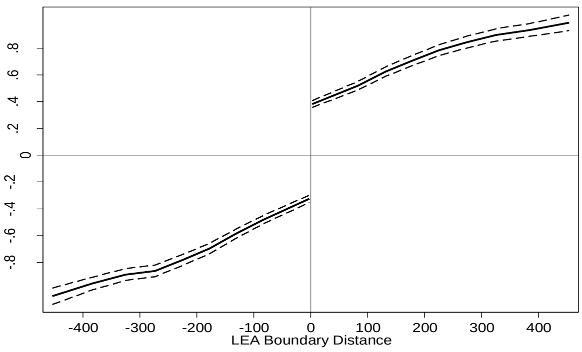

Log house price, by autonomous value-added, p=0.452

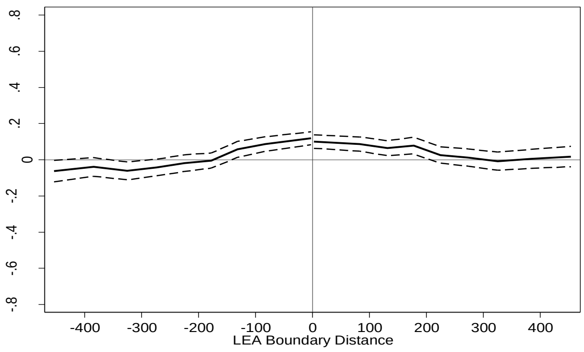

Notes: The scale on the x-axis is in metres from the boundary, at the minimum of each bin used in the regressions. The scale on the y-axis is in standard deviations.

Appendix A: Procedure for defining school catchment areas

In our 'revealed preference' procedure, we start by estimating the approximate shape of the catchment area for each school using the residential postcode of pupils in the year when they start at the school. This shape is delineated by the 75^{th} percentile of the home-to-school distance in each of 10 sectors radiating from each school location (starting West and moving anticlockwise). Each of the 10 sectors is drawn to capture 10% of the school intake. This procedure relaxes constraints on the shape of catchment areas, allowing for geographically asymmetric patterns of attendance with sufficient flexibility to apply our boundary discontinuity design. The reason we truncate the catchment areas at the 75^{th} percentile home-school distance in each direction is to remove outliers that could artificially inflate the size of the imputed school catchment areas. Discarding these outliers reduces the likelihood of erroneously drawing catchment areas across LA boundaries, and ensures that we focus on areas in which there is a high chance of admission. Note that we experimented with other distance thresholds, as well as with overlapping fixed interval radial sectors and alternative starting points and orientations, with little effect on the results.

It is important to point out that this procedure is required to assign school quality to housing transactions *within* LAs in an institutional setting where school catchment areas are not formalised and enforced. However, this approach does not drive the boundary discontinuities that we exploit to identify the causal impact of school quality on house prices. These discontinuities are generated by LA admissions rules discussed in Section 3.2, which dictate that students attend a school in their LA of residence.

It is worth highlighting why this fairly complex shaping procedure is necessary to our cross-LA boundary design by considering some alternatives. Suppose we simply assigned the quality of the nearest school to each housing transaction, or arbitrarily drew a circular catchment area around each school. To implement a boundary discontinuity strategy, we would need to artificially impose the constraint that a student in a house on one side of an attendance district boundary – i.e. the LA boundary – cannot attend their nearest school if it lies on the other side. Without this restriction, the set of schools available close to an admissions zone boundary, but on opposite sides of it, would be nearly identical to each other and there would be no source of variation in school quality for identification in the boundary discontinuity model (violating Assumption A2). However, we would not want to impose this constraint if the discontinuity did not actually exist. Our imputation procedure does not force any such truncation of the catchment area at the boundary unless it is supported by the spatial distribution of pupils' homes in relation to the schools they attend. Stated differently, we allow our *de-facto* catchment areas of schools close to the LA boundaries to be truncated and shrunk in the direction of the boundaries – as well as in any other areas and trajectories – only when the data reveal that this is the 'right' pattern. Therefore, any discontinuity that we detect in our analysis is *not* imposed by design but revealed by parental responses to institutional features and LA boundaries, and expressed by actual school choices.

Appendix B: Benchmarking the price effects against private school fees and labour market returns to education

The price response for a standard deviation in the *pupil* score distribution (2.7 value-added points) is around 11% or about £20,500 at the house prices prevalent at the time of our study (or approximately £1500 per year on a repayment mortgage over 25 years at 5% interest rate). This cost is equivalent to just over 2.5 years of private schooling fees (about £2800 per term for private day-schooling in England in 2006-7). These figures are also comparable to the value of the investment in a child's education. To see this, consider that Machin and McNally (2008) estimate a labour market return of about 0.42% to a one percentile increase in age-10 test scores, for a cohort of children raised in the 1970s and 1980s. This implies that a one standard deviation improvement in achievement at this age raises future earnings by 12%. Following Machin et al. (2010), we calculate the present value of this 12% increase on earnings between ages 16 and 65, discounted back to child's age 5, when parents are likely to buy their home for primary school admission.[13] This calculation gives a discounted lifetime benefit of approximately £20,600, which is very close to the house price response to one standard deviation improvement in the pupil test score distribution (about £20,500). This comparison is based on the fully capitalized value of the house, and the benefits of this investment could clearly outstrip the user costs taking into account potential house price appreciation. Similarly, the benefits could significantly outweigh the costs for families with more than one child. Nevertheless, these calculations illustrate that the house price response to school quality that we document is of a plausible magnitude given the expected return in terms of future earnings.

[13] Machin et al. (2010) estimate average yearly earnings for all individuals aged 16 to 64 in the Family Earnings Survey (2002/2003) to be around £10,700. They then propose to use a discount rate of 3.5%, in line with the recommendations in the UK HM Treasury Green Book (http://www.hm-treasury.gov.uk/data_greenbook_index.htm). Considering the 12% return to a one percentile increase in age 10 test scores discussed above, we estimate the benefits over ages 16 to 65, and discounted back to age 5, as follows: $NPV = \frac{1}{(1+3.5\%)^{11}} \sum_{t=1}^{50} \frac{£10{,}700}{(1+3.5\%)^{t}} \times 12\%$.