Spatial Economics Research Centre

SERC DISCUSSION PAPER 18

# Valuing School Quality Using Boundary Discontinuities

Stephen Gibbons (SERC, CEP and LSE)
Stephen Machin (CEE, CEP and LSE)
Olmo Silva (SERC, CEP and LSE)

March 2009

This work was part of the research programme of the independent UK Spatial Economics Research Centre funded by the Economic and Social Research Council (ESRC), Department for Business, Enterprise and Regulatory Reform (BERR), the Department for Communities and Local Government (CLG), and the Welsh Assembly Government. The support of the funders is acknowledged. The views expressed are those of the authors and do not represent the views of the funders.

# Valuing School Quality Using Boundary Discontinuities

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

**March 2009**

* SERC, Department of Geography and Environment, LSE and CEP
** Department of Economics, UCL, CEE and CEP
*** SERC, Department of Geography and Environment, LSE and CEP

Acknowledgements
We would like to thank Amy Challen and Anushri Bansal for excellent research assistance, and Jaap Abbring, Victor Lavy, Erik Sørensen and participants at the Workshop on Residential Sprawl and Segregation 2007 (Dijon), the CEP Annual Meeting 2008 (Cambridge), the CEPR Economics of Education and Education Policy in Europe Conference 2008 (Amsterdam), and seminars at CPB - The Hague, Sussex University and Tinbergen Institute for comments and suggestions. We are responsible for any errors or omissions.

Abstract
A large body of international research shows that house prices respond to local school quality as measured by average test scores. But better test scores could signal better expected academic outputs or simply reflect higher ability intakes, and existing studies rarely differentiate between these two channels. In our research, we simultaneously estimate the response of prices to school 'value-added' and school composition to show more clearly what drives parental demand for schools. To achieve consistent estimates, we push to the limit the use of geographical boundary discontinuities in hedonic models by matching identical properties across admissions authority boundaries; by allowing for a variety of boundary effects and spatial trends; by re-weighting our data to only consider the transactions that are closest to education district boundaries; and by submitting the estimates to a number of potentially destructive falsification tests. Our results survive this battery of experiments and show that a one-standard deviation change in either school 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 divided and ambiguous when it comes to answers about what makes a school a 'good' one or about what people value in education. Although parents making school choices seem well aware of their personal preferences and go to great lengths to secure a place for their children at their preferred schools, social scientists have had mixed success in eliciting any general conclusions about these preferences.

Researchers in education and the sociology of 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 suggests that, although parents rank academic outcomes highly among the reasons for choosing a school, there are many other factors that play an important role, such as distance from home, school composition and the child's potential wellbeing at school. More recently, actual choices regarding schools and teachers expressed by parents have been used as an alternative way to uncover parental preferences for school attributes. For example, Hastings et al. (2005) use parents' ranking of preferences in the US Mecklenburg County school choice program in North Carolina to document the value of schools close to home and of schools with high average test scores. In contrast, Jacob and Lefgren (2007) use parental requests for specific teachers to show that parents strongly prefer primary school teachers who are good at promoting student satisfaction, while they place relatively less value on a teacher's ability to raise standardised test scores.

These works aside, by far the vast majority of research in the field has looked for evidence of the value of schools in the capitalisation of their benefits into housing prices – i.e. 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) provide a summary of recent evidence, suggesting a consensus estimate of around 3-4% house price premium for one standard deviation increase in school average test scores. Bayer et al. (2007) offer a structural modification based on discrete housing choices that provides a correction to the standard hedonic framework when preferences are heterogeneous, and come to similar conclusions.

A limitation of this line of work is that – with only a few exceptions – it is confined to showing that prices tend to be related to headline school performance measures based on school average test scores. But better school average test scores could occur through improvements in school intake or through faster pupil progress within schools – potentially driven by teaching quality or resources. Existing studies rarely differentiate between these channels of influence. In fact, one possibility is that parents pay for school output or value-added because it represents what they expect their children to *gain* academically. A second possibility is that parents pay for good peers and favourable school composition – which are school inputs – irrespective of the likely contribution that these factors make 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).

Despite these compelling issues, the hedonic literature usually pays only passing attention to anything that parents might value in a school apart from pupils' final achievements. One exception is the work by Brasington and Haurin (2006), whose results appear to show that that school value-added and initial achievements both have positive effects on prices, although this important point is somewhat lost in their conclusion that value-added does not matter. Kane et al. (2005) also consider value-added and average test scores as alternative indicators of school performance. However, they do not present specifications in which both indicators are included at the same time, and do not provide (nor aim to provide) persuasive evidence on the importance of value-added . On the other hand, Clapp et al. (2007) show that the demographic characteristics (specifically ethnicity) of school pupils (i.e. inputs) seem more important than test scores (i.e. outputs) to home buyers around Connecticut schools, although the authors

[1] See Kramarz et al. (2009) for a detailed discussion, together with empirical tests, of the relative importance of pupil, school and peer effects in determining test scores.

do not have access to data on pupils' academic progress. Other papers have looked at the importance of school expenditure (and other inputs) 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. Very recently, Cellini et al. (2008) use referenda outcomes in California's school finance system to suggest that house prices respond to 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 sate 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. Gibbons and Machin (2006) suggest that popularity in itself raises prices, given that over-capacity schools command an additional premium relative to under-capacity schools with equal performance.

Clearly, from an educational policy perspective and for those interested in the processes of residential sorting, it matters which of these various drivers – school effectiveness, school composition or school resources – is important. This is because housing price patterns could reveal preferences for dimensions of school quality that are open to policy intervention, but are currently neglected by the narrow focus on test scores. In our paper, we take on some these challenges – mainly by differentiating between the impact of school value-added and school composition – to show that value-added is an important factor behind the house price response to school  test scores.

To carry out this analysis effectively we need reliable methods that take into account potential omitted variables and endogeneity issues (such as neighbourhood amenities). To achieve this, we apply, improve and test the boundary discontinuity regression methods that have become favoured in the field. Previous examples in the education research domain include Black (1999), Kane et al. (2005) and Fack and Grenet (2008), who analyse the impact of school test scores on house prices, and Bayer and McMillan (2005), in the context of school choice. Closely related thinking provides the foundation of studies that investigate the effects of market access when there are changes in national borders or their permeability. Examples include Redding and Sturm (2008), who look at changes that occurred during German division and re-unification, and Hanson (2003) who focuses on the opening of Mexican border as a result of the North American Free Trade Agreement. In a similar vein, boundary discontinuities have been used to assess the effect of taxation on housing prices (Cushing, 1984) and on the location and growth of manufacturing firms (Duranton et al., 2006).

One important innovation in our paper is to extend the boundary discontinuity methods to a context in which school admission zones are fuzzy, overlapping and only partially bounded. Additionally, we push the use of geographical boundary discontinuities in hedonic models to the limit by carrying out a range of novel robustness and falsification tests. These tests go much further than previous work in demonstrating that the housing prices are causally related to school attributes, and not spuriously correlated because of unobservable cross-boundary trends in neighbourhood amenities. Finally, one further contribution of our work is to apply these methods to the population of housing transactions in England and to schools data for the entire national system, thus improving the general validity of our findings.

To preview our methods, we first of all match properties with observably identical characteristics, within the shortest possible distance across education admissions authority boundaries, and estimate our models using differences between these cross-boundary pairs. Next, we weight our estimates towards the closest spaced sales pairs and control for a variety of boundary effects and spatial trends in prices across boundaries. We then go on to demonstrate the robustness of our results and the credibility of the boundary discontinuity approach by showing that prices are influenced only by the characteristics of schools that admit pupils on the basis of where they live: house prices do not respond to the quality of local schools that predominantly admit pupils on the basis of their religious affiliation, irrespective of where they live. Further, we check that there are no school-related price differentials for schools of different quality, but within the same school admission district, or across imaginary boundaries that do not delineate actual school admissions district zones (i.e. we fictitiously shift boundaries by 10km to the

North and 10km to the East). Our results survive this battery of experiments and falsification tests, and show that a one-standard deviation change in final test scores, brought about either school value-added or prior achievement, raises prices by around 3%.

The remainder of the paper has the following structure. Section 2 explains our methods. Section 3 discusses the context in which we apply these methods and the data setup. Section 4 presents our results and discussion, focussing firstly on identification of the effects of school performance on house prices, and then considering the role of value-added and school composition in this relationship. Finally, Section 5 provides some concluding discussions.

## 2. Empirical strategy

### *2.1. Methodological framework*

Our empirical work uses a regression discontinuity design that builds on the geographical 'boundary discontinuity' approach. This method was popularised for use in property value analysis by the work of Black (1999) and has been employed several times since (e.g. Bogart and Cromwell, 2000; Gibbons and Machin, 2003, 2006; Kane et al., 2005; Davidoff and Leigh, 2007; Fack and Grenet, 2008; Bayer et al., 2007). The standard 'hedonic' property value model is well known (Sheppard, 1999) and tries to explain property values (or, most commonly, log property values) as a linear combination of observable property attributes and the 'implicit prices' of these attributes in the housing market. These implicit prices can be estimated by standard least squares regression techniques, but the pervasive problem with this approach is that researchers do not observe all salient property and neighbourhood characteristics, leading to serious omitted variable issues. This problem is particularly acute when neighbourhood amenity quality and local public good quality – like school quality – depends on the distribution of characteristics in the local population. In such cases, any unobserved attribute that raises local housing prices changes amenity quality through residential sorting, because higher price houses are (on average) occupied by higher income households.

One way to mitigate this problem is to compare only close-neighbouring houses, because these often tend to be quite structurally similar and self-evidently have identical, or very similar, neighbourhood environments. So one can potentially eliminate area effects in a house price model by taking differences between houses that are in close proximity or modelling neighbourhood fixed effects. However, this strategy is not useful if the goal of the research is to obtain implicit prices of neighbourhood attributes, local amenities or public goods, unless there is a sharp discontinuity in the supply of these attributes between close-neighbouring homes.

This last condition holds when school admissions are arranged according to contiguous pre-defined admission zones: residents on one side of the boundary have access to a different school or set of schools than do residents on the opposite side of the boundary. A researcher looking at the effect of schools on house prices can therefore mitigate the problems caused by unobserved neighbourhood attributes by including attendance district boundary dummy variables in regression models (unless the boundaries are particularly long), or by working with differenced data from a matched pair of neighbouring houses on either side of the boundary. The empirical model underlying this approach is set out below in a way that will help explain our empirical methods.

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

$$p = s(c)\beta + x(c)\gamma + g(c) + \varepsilon \qquad (1)$$

Where $s(c)$ represents the school 'quality' that home buyers expect to be able to access by residence at $c$, prior to school admission, measured on the basis of school characteristics at periods prior to the house sale. These attributes can be generally thought of as measures of school composition, resources and effectiveness. In our empirical application we will try to estimate the effects of these different components separately. As usual, $\varepsilon$ 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 spatial locations, such that the price varies with geographical location, for example due to unobserved neighbourhood characteristics and amenities (other than schooling).

Location $c$ can be specified in various ways, most flexibly in terms of a vector of geographical or Cartesian coordinates. We discuss this in more detail below.

*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 the unobserved attributes of location $c$. A spatial differencing strategy offers one way to eliminate common area effects $g(c)$. Taking differences between specific houses $i$ and $j$ results in the following specification:

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

This transformation, at least on its own, does not appear to offer advantages. Least squares estimates of the implicit prices ($\beta, \gamma$) are consistent if and only if the difference in unobservable price determinants $g(c_i) - g(c_j)$ is uncorrelated with the difference in school quality $s(c_i) - s(c_j)$ and with differences in other housing attributes $x_i(c_i) - x_j(c_j)$. This condition will not hold in general, and consistent estimation of $\beta$ requires the researcher to find locations $i, j$ such that locally $Cov[s(c_i) - s(c_j), g(c_i) - g(c_j)] = 0$ and $Var[s(c_i) - s(c_j)] \neq 0$ (conditional on observed housing and neighbourhood characteristics). These two conditions will never be met simultaneously and exactly, except for pathological cases[2], for any continuous functions $s(.), g(.)$ because the first condition 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:

[2] For example if $\frac{\partial s(c)}{\partial c} = 0$, or $\left.\frac{\partial s(c)}{\partial c}\right|_i = \left.\frac{\partial s(c)}{\partial c}\right|_j$ and $\left.\frac{\partial g(c)}{\partial c}\right|_i = -\left.\frac{\partial g(c)}{\partial c}\right|_j$ such that $Cov[s(c_i) - s(c_j), g(c_i) - g(c_j)] = 0$.

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 $\theta$ is a positive constant (or positive definite matrix if $s$ is multidimensional).[3]

The geographical 'boundary discontinuity' approach effectively amounts to an attempt to exploit 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 standard regression discontinuity designs (Imbens and Lemieux, 2008) in which a single forcing variable (e.g. voting share, such as in Lee et al., 2004) determines 'treatment' (e.g. party affiliation of elected representative), although the general principle is similar.

In practical empirical settings, there are three main reasons why the identification strategy sketched above could fail:[4]

(a) There are spatial trends in amenities across boundaries such 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.

(b) There are boundary discontinuities in prices, not caused by school quality differences, which violates assumption A1.

(c) School quality lacks any discontinuity at attendance boundaries, violating assumption A2.

---

[3] Note that assumption A2 is a necessary condition if there is to be any variation in school quality to allow estimation of an associated hedonic price. On the other hand, A1 is sufficient, but not necessary, given the pathological cases outlined in footnote 2.

[4] One additional assumption is that $g(c)$ represents a spatially isotropic process, so that direction does not matter and buyers do not care more about, say, bad neighbours to the left than bad neighbours to the right. If this is not the case then even identical co-located properties may have different prices depending on which way buyers are looking when they make their valuation.

Regarding case (a), highly localised factors (e.g. a noisy next door neighbour) that influence sales prices of individual homes, but are uncorrelated over space (i.e. they are 'noise', contained in $\varepsilon_i - \varepsilon_j$ ) are not of serious concern. These property-specific factors do not affect housing market prices in a way that could influence school quality through population sorting. However, we *do* need to be concerned about spatially correlated amenities and other desirable attributes that could lead house prices on one side of a boundary to differ *on average* from house prices on the other side. This situation could arise if, for example, one attendance zone contained a rail station and another did not (see Gibbons and Machin, 2005, for evidence of the amenity value of rail access). This would then result in higher prices, richer families and better schools in the 'station zone', and a spatial trend in house prices rising across the boundary towards the station. Because of this trend, the price differential between houses on different sides of the boundary grows with the distance between sales. Hence we could find a correlation between house prices and school quality amongst closely spaced neighbours that is not caused by the demand for school quality, but by residential sorting that is a consequence of demand for rail access.

Even if there are no gradual cross-boundary price trends, there can be cases of type (b), where prices change sharply from one side of the boundary to the other. First, administrative attendance zone boundaries may coincide with distinct geographical features, e.g. major roads, which partition communities. Then, if these communities are different, the boundary may create a discontinuity in average housing prices over short distances that is not school-related, violating the assumption that $g(c)$ is continuous. Secondly, even without visible evidence of the boundary on the ground, houses on different sides of a boundary could have different directional aspect or outlook. Consider, for example, two long rows of houses on an east-west running boundary, one with sunny gardens facing south and one with shady gardens facing north. If residents with children prefer sunny gardens, then this aspect could be sufficient to induce a housing price differential and a consequent school quality difference across the boundary. Thirdly, contiguous districts may have different tax rates or offer different district-specific amenities, generating a sharp discontinuity in prices that is not caused by schools.

Lastly, lack of discontinuity of type (c) occurs if attendance boundaries do not, in practice, act as a barrier to pupils attending schools in districts neighbouring their homes. This could happen, for example, if changes in school policy have removed the importance of traditional attendance zones. Note however, that even if some pupils can cross these boundaries, condition A2 will still hold. In fact, identification (in the sense of condition A2) requires only that there is a discrete jump in the probability of attending schools on different sides of the boundary as one moves from a residence on one side to a residence on the other, but this change in probability need not be from zero to one – i.e. the discontinuity can be fuzzy (Imbens and Lemieux 2007). This change in probabilities ensures that there is a discrete jump in *expected* school quality (before admission) from one side to the other.

### *2.3. Proposed methods to address the identification problems*

A few of these identification concerns have been considered and partly addressed in the earlier literature. However, in this paper, we take these problems into much deeper consideration and go a long way further than existing work in establishing the credibility of the boundary discontinuity approach in our empirical context. With this purpose, we extend the standard methodology and produce a series of powerful robustness and 'falsification' checks. These key extensions and tests are as follows (numbered method M1-M8 for recognition in the Results section below):

M1. *Visually assess and (statistically) test for the presence of discontinuities*: Drawing on the regression discontinuity design literature (and similar to Bayer et al., 2007 and Kane et al., 2005), we provide some graphical evidence and statistical tests regarding such discontinuities in area characteristics.

M2. *Match property transactions with identical observable characteristics across administrative boundaries*. We pair up each house sale with the nearest transaction on the opposite side of an administrative attendance district, where the transaction is of the same property type and occurs in the same year (see also Gibbons and Machin, 2006, and, to a lesser exten,t Fack and Grenet 2008). This approach borrows from the literature on discrete-cell matching, first pioneered by Rubin

(1973), which accounts for the effect of observable characteristics on the outcome of interest in a fully non-parametric way. In our set-up, this equates to allowing the price effects of matched property characteristics to vary by boundary.

M3. *Weight regressions to zero-distance housing transaction pairs.* Earlier work (e.g. Black, 1999) tested robustness to cross-boundary trends by selecting houses in increasingly narrow distance bands along either side of the boundary, that is applying weights of 1 to transactions within a specified boundary distance, and weights of 0 to those outside that distance. We extend and generalise this idea 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. Re-weighting our analysis in this way ensures that our identification predominantly comes from observations where the identifying assumptions A1 and A2 are most likely to hold.

M4. *Include boundary fixed effects in cross-boundary difference models*. Our institutional context (described below in Section 3) offers us multiple schools on each side of an attendance district boundary, so school quality varies across boundaries *and* along a boundary within a given attendance district. This data structure means we can control for boundary fixed effects (using boundary dummy variables) in our cross-boundary differenced model, thus eliminating between-boundary variation. This is crucial given assumption A1 and the problems with boundary-specific discontinuities highlighted in Section 2.2 under case (b).

M5. *Control for distance-to-boundary trends and polynomials*. We follow the regression discontinuity design literature by controlling for polynomial trends in 'distance' from the discontinuity (e.g. DiNardo and Lee, 2004; Lee et al., 2004; and Clark, 2008). In our context, this 'distance' is literally the geographical distance from attendance district boundaries. Like other studies in this field, we impose some parametric structure, for example by specifying that

$g(c_i) - g(c_j) = \rho_{11}d_i + \rho_{12}d_i^2 + \rho_{13}d_i^3 + \rho_{21}d_j + \rho_{22}d_j^2 + \rho_{23}d_j^3$, where $d_i$ is the distance from sale $i$ to the boundary, and $d_j$ is the distance from the matched sale $j$ to the boundary. Note that we can further control for different trends for each boundary, by including boundary dummy × distance to boundary polynomial trends; and for asymmetric trends on opposite sides of boundaries, for example by interacting distance polynomials with an indicator of whether the school admission district is 'better' or 'worse' in terms of school quality than the adjacent one. The latter experiment mimics and generalises the approach taken by the regression discontinuity literature, which interacts distance polynomials with indicators identifying whether observations are to the 'left' or to the 'right' of the discontinuity point (e.g. treated or controls in the case of binary treatments). By explicitly modelling trends in prices as we move away from school district boundaries we act to mitigate the issues discussed under point (a) in Section 2.2.

M6. *Restrict our attention to boundaries where pupils rarely cross*. Our data uniquely allows us to observe when pupils cross an admission district boundary to attend their school. Thus, we can check that our results are not compromised by the 'fuzziness' of the school quality discontinuity, or by the lack of it, caused by excessive pupil movements across boundaries. This allays the concerns highlighted in point (c) in Section 2.2.

M7. *Apply falsification tests using 'fake' attendance boundaries*. We re-estimate our models using differences between transactions in the same attendance district and using differences between property transactions along imaginary attendance boundaries, created by translation of the geographical coordinates. While the first method was applied in Black (1999) , the use of completely artificially translated boundaries is novel and provides a powerful and stringent falsification test. A finding of a positive association between school quality and housing prices in this setting would falsify the claim that price effects are causally linked to cross-boundary school quality discontinuities. This exercise thus helps to allay some of the concerns raised in point (a) in Section 2.2, and helps to verify the validity of assumption A1.

M8. *Compare the methodology for cases in which home location is and is not a school admission criterion*. Our institutional context provides us with two types of schools. For 'non-autonomous' institutions, places are allocated by the admissions authority predominantly according to how close a pupil lives to the school, and attendance district boundaries are binding.[5] There are therefore strong reasons to buy a home close to a school of choice, and on the 'right' side of the boundary. On the other hand, 'autonomous' schools – mainly religious – operate pupil admissions policy autonomously of local authority control and families need *not* buy their home close to the school to gain admission. In fact, for religious schools, regular attendance at designated churches and other expressions of religious commitment are foremost. Although parents might still want to reside close for convenience and to minimise travel costs, they do *not need* to do so to secure admission to their children. Thus, local house prices will not respond to the quality of 'autonomous' schools since these can be accessed irrespective of residence. These institutional features provide us with an opportunity to run a particularly demanding 'falsification' test. By comparing the response of house prices to both types of schools, we can ascertain whether any links between school quality and house prices are genuinely attributable to the demand for those school characteristics across boundaries, or else spurious. This provides an additional tool to check the issues raised in points (a) and (b) in Section 2.2.

The robustness and falsification tests described above relate to identification of the causal effect of school quality and other characteristics on house prices. We now turn to describe an additional set of identification issues that arise when the research goal is to interpret the above estimates as 'willingness to pay' for school quality.

[5] We will discuss the features of the school admission system in England that are most relevant to our research in Section 3.2. Gibbons et al. (2008) deal with this issue in much greater detail.

### 2.4. *Identification in hedonic models when there is sorting on school quality*

It is well known that empirical identification of marginal willingness to pay for any neighbourhood amenity in a hedonic model is challenging when different households have different incomes and different preferences for this amenity, leading to residential sorting. Under these conditions, the distribution of household characteristics near good quality schools will be different from the distribution of characteristics of residents near poor quality schools, even if school quality is the only factor determining house prices. This sorting has two consequences.

Firstly, linear regression estimates may not provide estimates of the *mean* valuation of school quality, because the marginal willingness to pay (WTP) for school quality varies across the distribution of household characteristics. Allowing for heterogeneity by interacting school quality with household characteristics (e.g. income) is a poor solution, because if WTP varies by characteristics, then these characteristics are endogenous in house price regression models. Bayer et al. (2007) focus on this particular identification problem and describe a solution using a two-stage structural approach that imposes a particular functional form on the residential choice and sorting process (coupled with an instrumentation strategy). In particular, the authors follow the approach of Berry et al. (1995) and first specify a functional form for the indirect utility function of a household with given set of characteristics and given housing choice. This depends linearly on the characteristics of the housing choice and of the surrounding neighbourhood, plus interactions between these attributes and household characteristics. Then, they go on to estimate a multinomial logit model on actual housing choices to retrieve the set of parameters that characterise the mean indirect utility function of all households in a given housing choice, and the household specific components. Finally, in the last step of their procedure, the authors use the estimated parameters to control for the effect of heterogeneous preferences in a standard hedonic price regression.

Although technically impressive, this method relies on strong and hard-to-test assumptions about the shape of the indirect utility function and on the Independence of Irrelevant Alternatives (IIA) hypothesis invoked to estimate multinomial logit models. It is thus difficult to generalise its applicability and understand the consequences of the failure of any of the required assumptions. In our work, we do not wish to impose this much structure, but present no novel solution to these issues. In the presence of heterogeneous preferences and/or incomes and sorting across boundaries, our discontinuity design will provide a weighted average of the marginal WTP of residents along the admissions zone boundary. This estimate may be an upward or downward biased estimate of *mean* marginal WTP. However, in our defence, the work by Bayer et al. (2007) shows that, both empirically and from a theoretical point of view, the 'traditional' hedonic models *are* effective at evaluating mean WTP in contexts (like ours) where the amenity in question is supplied at various qualities in many different locations.

For the same reasons, in this paper we also do not consider the issue of heterogeneity in the responses of house prices to school quality depending on buyers' or neighbourhood characteristics. These are endogenous in house price regression models in the presence of sorting, and cannot be simply added to empirical specifications in interaction with school quality.

The second consequence of sorting on school quality is that it makes it difficult to separately identify the marginal willingness to pay for school quality from the willingness to pay for neighbours' quality. Neighbourhoods with access to better schools will usually be home to those households with higher income and with stronger preferences for school quality. Hence part (though clearly not all) of the association of between school quality and house prices works through its effect on neighbour quality, so estimates cannot be easily interpreted as WTP for school quality *per se*. 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, we will demonstrate that our estimates of the value of school quality are steadfastly linked directly to school attributes, and in this control function context not to neighbourhood quality.

## 3. Institutional context and data setup

Before presenting our results in the next Section of the paper, we first offer a description of England's primary schooling system in more detail. We also discuss the data sources that we use to implement our work and the empirical specifications that we consider.

### *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.[6] 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 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 are intended to 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, 2009).

In our empirical work below, we use the *ks1* to *ks2* value-added score (*va*) as the main indicator of schools' production *output,* or effectiveness. On the other hand, we treat *ks1* scores as a general control

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

for pupils' prior academic achievements, i.e. mainly as a measure of school *inputs* in terms of the educational advantages embodied in the composition of its pupil intake. These *ks1* tests might, in part, reflect the effectiveness of a school in children's early years. However, they are not publicly available and so cannot provide parents with a direct signal of school performance. Thus, we treat *ks1* scores as capturing information about school composition that parents can only learn about from school visits, word of mouth, and using local knowledge. Our results in the following sections seem to confirm that *ks1* test scores are predominantly linked to students' background characteristics. Note additionally that our main justification for focusing on *ks1* scores as an indicator of background (rather than income or free school meal eligibility) is that the coefficient on value-added conditional on *ks1* in our regressions can be easily interpreted in terms of pupil progress *or* final achievement.

### *3.2. School types and admissions*

All state primary schools in England are funded largely by central government, through Local Authorities (LAs, formerly Local Education Authorities) that are responsible for schools in their geographical domain.[7] These schools fall into a number of different categories, and differ in terms of the way they are governed and who controls pupil admissions.[8] Most primary schools (roughly two-thirds) are termed 'Community' schools and are closely controlled by the LA. Other types of school, instead, are usually linked to a Faith or other charitable organisation, and more autonomously run. The key difference relevant to this paper is between schools that administer their own admissions and 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) provide more details on the overall differences between these two groups of schools.

---

[7] Broadly speaking, LAs are responsible for the strategic management of education services, including planning the supply of school places, intervening where a school is failing and allocating central funding to schools.

[8] In addition there is a small private, fee-paying sector, which we do not consider here. Private schools educate around 6-7% of pupils in England as a whole

Regarding pupil admissions, overall, all LAs and schools must organise their 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 designated churches and other expressions of religious commitment are of foremost importance. Place of residence, in contrast, almost never features as a criterion. Even then, if place of residence is important for admission, it relates to Diocese boundaries, which do not follow administrative and 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.

There is however one additional crucial feature of the admission system that applies to non-autonomous, but not to autonomous schools, that we exploit in our empirical work. Pupils rarely attend non-autonomous schools outside of their LA of residence. Families are allowed to apply to non-autonomous schools in other LAs, but up until recently (covering the period we consider in our empirical work) parents had to make separate applications to different LAs. More importantly, LAs do not have a statutory requirement to find a school for pupils from other school districts: the law only requires that they provide enough schools for pupils in "their area".[9] As a result, banking on admission to a popular non-autonomous school in another LA is a high-risk strategy and LA boundaries act as admissions district boundaries over the period we study. This provides a source of discontinuity in the non-autonomous school 'quality' that residents can access on different sides of LA boundaries. In contrast,

[9] 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".

these barriers are much less relevant for admission to Faith schools and other autonomous schools that manage their own admissions. In Section 4.2 below, we will provide clear and compelling evidence that LA boundaries significantly affect non-autonomous school attendance patterns, and that there is a discrete jump in the probability of attending schools in a given admission district as one moves from a residence on one side to a residence on the other side of a boundary.

### *3.3. Source data*

In our analysis we combine information obtained from three different data source. First, we use data from the UK Land Registry. This is an administrative dataset that records the address, sales price and basic characteristics (property type, new or old build, free-hold or leasehold) of all domestic properties sold in the UK. This "Price-paid" dataset is available from the year 2000 onwards and provides the housing sales price information which is central to our research. Each property is located by its address postcode, typically 10-12 neighbouring addresses, and each postcode can be assigned to a 1 metre coordinate on the British National Grid system using the National Statistics Postcode Directory.

Next, for information on school quality and characteristics, we rely on data from the UK's Department for Children, Schools and Families (DCSF ). The DCSF collects a variety of census data on state-school pupils centrally, because the pupil assessment system is used to publish school performance tables and because information on pupil numbers and characteristics is necessary for administrative purposes – in particular to determine funding. A National Pupil Database exists since 1996 holding information on each pupil's assessment record in the Key Stage Assessments throughout their school career. Since 2002, a Pupil Level Annual Census (PLASC) records information on pupil's school, gender, age, ethnicity, language skills, any special educational needs or disabilities, entitlement to free school meals and various other pieces of information including postcode of residence. PLASC is integrated with the pupil's assessment record in the National Pupil Database (NPD), giving a large and detailed dataset on pupils along with their test histories. Additional institutional characteristics and expenditure information on schools is obtained from “Edubase” data, from the Annual School Census and from the Consistent Financial Reporting series that can be obtained from the DCSF.

Finally, neighbourhood characteristics from the 2001 GB Census at Output Area level can easily be linked to the Price-paid housing transactions data by their address postcode. We can also compute various geographical attributes such as distances to LA boundaries and distances between properties using a Geographical Information System. Linking the schools data to housing sales is however more complex, since there is no predefined mapping between a house sale, i.e. its postcode, and the set of schools that are accessible from that location. We infer this mapping from the data available to us and rely on actual home-school travel-patterns to draw school catchment areas based on revealed preferences. We describe this computationally intensive, but intuitively simple procedure in the next section.

*3.4. Linking schools to housing transactions and matching across boundaries*

One of the innovations in this work is the accurate assignment of school quality to house location in our data, when the institutional context means that there is not a one-to-one mapping between where a child lives and the school he or she attends. The procedure entails imputation of the set of schools accessible from each postcode in our Land Registry housing transactions database using the attendance patterns of pupils as recorded in our National Pupil Database. This approach is more sophisticated than those previously adopted when school admission zones are fuzzy, overlapping and only partially bounded (like in the UK) – e.g. assigning house to the nearest school or set of schools. Defining catchment areas from ‘revealed preferences’ in this way implicitly accounts for features of school choice and attendance patterns that would be obscured by more restrictive assumptions, such as common travel distances or common travel modes in different urban and rural contexts.

In our revealed preference procedure, we start by estimating the approximate shape of the ‘catchment’ area for each school using the residential addresses (postcode) of pupils in the year when they start at the school (provided in the integrated PLASC-NPD). 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. Each of the 10 sectors is drawn to capture 10% of the school's 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. Note that we truncate the catchment area at the 75$^{th}$ percentile home-school distance in each direction in order to remove outliers could that artificially inflate the size of the imputed school catchment areas. We have experimented with higher percentiles of the distance distribution to delineate the limits of the catchment area, and have used alternative limits such as four times the median distance in each direction. However, drawing the limits close to the most distant pupils creates implausible, large areas, which imply that houses have access to far too many schools over too wide a distance. Evidently, the addresses of pupils who live a long way away from the schools they attend, relative to their peers, provide poor information about the limits of school catchment areas and the real chances of admission. These outliers can come about through address coding errors, or because families gain admission to a school for their first child and then move away from a school before a subsequent child is admitted under a siblings rule. In conclusion, removing these outliers reduces the likelihood that we erroneously draw catchment areas across LA boundaries and ensures that we focus on areas in which there is a realistically high chance of admission – a consideration which is paramount to home buyers seeking to get their children into a particular school (and thus to our research).

Note that we have also tried using 20 overlapping fixed interval ($\pi/10$ radians) radial sectors to define the direction bands and shape of the catchment area, but this does not substantively change the findings that we discuss here. Moreover, in the results that we report below, we start from the West direction (relative to the centroid of the school postcode) when calculating the 10 radial sectors originating from the school, and then proceeded anticlockwise. We have experimented with alternative starting points and orientations, with little effect on our estimates.

Before moving on, let us emphasise the reason why this shaping procedure is necessary. This is best seen 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. We would then need to artificially impose the constraint that a student in a house on one side of an administrative attendance district boundary (i.e. the LA boundary) can not attend their nearest school if it lies on the other side. In fact, without this restriction, the set of schools available close to an admissions zone boundary, but on opposite sides of it, would be similar or identical to each other. Hence, there would be no source of variation in school quality for identification in the boundary discontinuity model (violating Assumption A2). On the other hand, if we imposed this constraint, we would enforce a discontinuity, although this discontinuity might not actually exist.

Crucially, the existence of these discontinuities is something that we do not want to assume or take for granted, rather a feature that we intend to test in our data. Our imputation procedure is flexible enough to allow us to do this. This is because it allows the catchment area of schools close to the LA boundaries to be truncated and shrunk in the direction of the boundaries, but not in other areas and trajectories. Once again, our procedure does not impose this truncation unless it is supported by the spatial distribution of pupil homes in relation the schools they attend.

After creating each school-specific catchment area definition (based on the residential postcode of pupils in PLASC at the time when they start at the school), we calculate the distance and direction from each school to each housing transaction in our Land Registry housing transactions database (up to a maximum distance of 10km). It is then straightforward (though time consuming) to link each house to multiple schools by deducing which housing transactions lie within which school catchment areas. Following that, we calculate variables summarising the set of schools that are accessible from a given housing transaction postcode in a given year, by aggregating the characteristics of the schools to which a house is linked. When we aggregate school characteristics, we apply a higher weight to the closest schools to each house, by inverse-distance weighting, although the results we present below turn out to be generally insensitive to whether we use weighted or un-weighted aggregates. In carrying out this aggregation we maintain the distinction between autonomous and non-autonomous schools. So, for example, a housing transaction is assigned the mean value-added of local non-autonomous schools and the mean value-added of autonomous schools as separate variables.

We also take care to correctly organise the timing of events in our data. The pupil census in England occurs in January, 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 test results and census figures published at the end of year *t-1* (in October to November).

The procedure described above yields a large dataset of over 1.6 million housing sales for 2003, 2004, 2005 and 2006 joined to data on the average characteristics of the set of schools that can be accessed from the postcode of each sale. To set up the spatially differenced cross-boundary model in Equation (2) we reduce our sample to the set of sales occurring within 2500m of a LA (attendance district) boundary. We then find, for each transaction, the nearest sale in the same year of the same property type, occurring in an adjacent LA, within the median inter-property distance across that specific boundary (method M1 in Section 2.3). This means that a given housing sale can provide a 'match' for multiple housing sales. Note that property type here is defined by detached, semi-detached, terraced or flats, and by ownership type, i.e. leasehold or freehold. Further, the restriction on matching within median distance along a boundary ensures that we do not create any matched pairs that are excessively far apart, given the density of houses in the local area. For reasons explained in Section 2.3, we also set up a set of matched sales across 'fake' LA boundaries and a set of matched sales within LAs (method M7). To produce the first sample, we simply translate the geographical coordinates of the housing transactions data by 10km North and 10km East, and repeat the matching exercise. For the second, we repeat the matching exercise but impose 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.

*3.5. Empirical specification*

Applying the data described above to the models of Equations (1) and (2) yields empirical specifications of the form:

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

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

where $p_{hi}$ is the (log) price of the house sale $h$ in location $i$. The variable $va_i$ represents the expected value-added, while $ks1_i$ represents the mean age-7 test score, for schools that can be accessed from location $i$, measured at periods prior to the house transaction. Vector $z_i$ represents other observable school and neighbourhood characteristics. Additionally, $x_h$ is a vector representing the observable attributes of house sale $h$. Finally, the function $g\left(c_i\right)$ represents unobserved neighbourhood characteristics and amenities (other than schooling) that might affect market prices. We parameterise $g\left(c_i\right)$ using boundary dummy variables, distance to school, distance between matched transactions and various distance-to-boundary polynomials. As usual, $\varepsilon_i$ represents unobserved housing attributes and errors that are independent of all other factors (i.e. 'noise'). The notation $\Delta$ means a difference between matched, closest transactions on either side of the LA boundary.

Note that we have sales and school attributes in multiple periods, but we have suppressed the *t*-subscripts for simplicity. Although variation over time in the cross-boundary differences in school quality contributes to identification, we do not exploit the time dimension alone in our estimation strategy. Three reasons for this decision are: a) test scores assigned to house postcodes are highly correlated from one period to the next so that the within-place, between-period variance in school quality is low; b) we have only 3 full years (2003, 2004 and 2005) and one quarter (quarter 1 of 2006) of housing transactions linked schools data; and c) response of prices to changes is likely to display inertia and be sluggish. These factors mean we cannot use changes over time alone as a basis for identification.

In the next section we present results from regression estimates of the models in (3) obtained by pooling all available time periods.

## 4. Results

### *4.1. Descriptive statistics*

Table 1 presents some key descriptive statistics. The first two columns summarise the full data set of housing transactions and associated school characteristics from 2003-2006. The second two columns present comparable statistics for our boundary sub-sample of sales, described in the Data section above. The average price of sales in the transactions data set is £182730. In the boundary sub-sample the mean is about £13000, or 7% higher. This is because administrative boundaries are more prevalent in and around towns and cities and hence we pick up more urban transactions in the boundary sub-sample. In addition, there is a greater chance of finding matched pairs of sales across sections of the boundaries in urban areas, where housing is more dense. It is easy to visualise this in Figure 1 and Figure 2, which plot the locations of transactions in the boundary sub-sample, for two arbitrarily chosen geographical areas: the Midlands, North West and South Yorkshire; and London and the South East. These figures illustrate a general spread of sales throughout England's cities and towns, but in a way that is governed by the administrative boundary structure.

In terms of school test scores, value-added is higher in the boundary sub-sample and *ks1* scores are lower, but the differences are relatively small. Houses in this sub-sample have slightly fewer accessible schools (where accessibility is imputed from travel patterns described in the Data section above). This difference is in accordance with our claim that LA boundaries restrict the choice set for houses located close to the boundary (see the discussion above and Gibbons et al., 2008). Schools also tend to be closer to home in the boundary sub-sample, again reflecting the relatively urban nature of the sample.

For the boundary group, we present some 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 is nearly 500 metres, and the raw average distance between matched properties is just under 725 metres. These figures look high in comparison with previous studies that focus on city neighbourhoods only, but are not so large in the light of the general geographical spread illustrated in Figure 1 and Figure 2. In our regressions, we apply inverse inter-sale distance weights, so the inverse distance weighted (IDW) means provide a better representation of the effective boundary difference relevant to our regressions. The effective mean distance to the boundary in the weighted sample is only 133m, and the weighted inter-sale distance only 206m.

### *4.2. Evaluating the boundary discontinuities*

As discussed at length in Section 2.2, a pre-requisite of our method is that a discontinuity exists in school quality at LA boundaries (or in the school quality households expect to be able to access; see Assumption A2). As a preliminary step, we show that, cross-district school attendance is much less prevalent than within-district attendance, even close to district boundaries. The relevant figures are presented in Table 2 and refer to proportions in the postcode. In the full dataset, only 3.3% of pupils attend schools other than in their home LAs, though this is not surprising given that, on average, schools in other 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 (given that the postcode has any children of primary school entry age) is 25%. Since this figure corresponds to addresses only 133m from the boundary (Table 1), we would expect nearly 50% chances of attending a school on either side of the boundary if it did not impose a 'barrier' and was unimportant for admission. Moreover, these means are from distributions that are highly right-skewed and the median proportion of pupils attending a school in a district different is zero. Clearly, then, LA boundaries create a strong impediment to school choice. This is fully consistent with the results using boundary discontinuities to identify the causal impact of school choice and competition on pupil achievement in Gibbons et al. (2008).[10]

[10] See also Card, Dooley and Payne (2008).

More explicit tests for discontinuities in school quality and other area characteristics at the LA boundary are provided in Figure 3 and Figure 4 (using method M2 of Section 2.3). In all these figures, the x-axis reports the distance from a property transaction 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). On the other hand, the left side of the diagram (distance < 0) corresponds to cases where access is to schools with value-added below that on the other side of the boundary. The plots are obtained as predictions (and with 95% confidence intervals plotted as dotted lines) from a regression of the standardised cross-boundary difference in the relevant variable, on a positive side and negative side constant term, and 18 distance-decile dummies, up to 800m from the boundary on each side.[11] The plots are restricted to 400m 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 absolute values of the positive and negative constants in the regressions are equal to one another). Note that the reason why these graphs are not necessarily symmetric is that a sale *i* on the 'good' side of the boundary may be matched with its closest sale *j* on the 'bad' side of the boundary, but sale *j* may in turn be matched to another sale *k* on the 'good' side of the (same or a different) boundary if *j* is closer to *k* than *i*. The standard errors are clustered on location $c_j$ to allow for repeated matches of the same sale *j* to multiple sales *i,* and for a degree of arbitrary spatial correlation in the error term.

The top left panel of Figure 3 shows a large and sharp discontinuity in value-added scores at LA boundaries for non-autonomous schools, making it clear that we have substantial variation in our main school performance measure across boundaries (Assumption A2). The overall scale of the difference within the 400 metres of the boundary is unsurprising given this is the variable on which the right and left halves of the plot are defined. However, the most important point here is that almost half of the 2-standard deviation spread occurs within the first 100m, from where our identification will predominantly

[11] The variables are standardised by the standard deviation of the cross-boundary difference within 800m.

come. The top right panel shows that a discontinuity in house sale prices exists too: although visually this looks small, the difference across the boundary is highly significant, and the price on the ‘good’ boundary side is higher than the price on the ‘bad’ boundary side at every corresponding distance. Rough visual comparison of the top left and right panels suggests that a 0.8 standard deviation change in value-added is associated with a 0.05 standard deviation change in house prices at the boundary. As we move away from the boundary, focussing on more widely spaced properties, we see that prices tend not to follow school value-added. This occurs because many other amenities drive these spatial price trends, illustrating the potential importance of weighting our regression estimates to close-neighbour observations, and controlling for distance-to-boundary trends (methods M3 and M5 in Section 2.3).

In the lower two panels of Figure 3, we look at the corresponding cross-boundary discontinuity picture for autonomous school quality. In these graphs, the right hand side corresponds to places with relatively high *autonomous* school quality (and vice versa for the left hand side). Again, there is by definition a strong rise in school quality across the boundary. However, there is no sizable discontinuity in house prices at the boundary in this institutional context, where admission to school is not so strongly linked to where pupils live. In fact the p-value of the F-test (= 0.76) shows that one cannot reject the null hypothesis of no cross-boundary difference in house prices.

In Figure 4 we present similar pictures for a whole range of neighbourhood-related characteristics, with left and right sides split by low and high non-autonomous school value-added. These plots serve to show to what extent cross-boundary neighbourhood differences are correlated with cross-boundary non-autonomous school value-added differences. It is evident that there are no discontinuities in terms of a wide range of neighbourhood characteristics (obtained from the 2001 GB census and the Land Registry data), including the share of local dwellings sold per year, the dwelling size and residents' characteristics. One exception is the proportion high-qualified residents (degrees and equivalent), in which there is a statistically significant break. The fact that more highly educated residents live on the side of the boundary with good schools is evidence for some degree sorting of those with higher incomes and stronger preferences for their children's education (similar results are found in Bayer at al., 2007). The empirical issues arising from this kind of sorting were discussed in 2.4, and we will address them in our regression robustness tests presented in Section 4.6.

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

Table 3 presents the coefficients and standard errors for our main regression results. We report only the key figures for the house price effects of school-mean value-added ('output') and *ks1* test scores, which we claim proxy for school 'inputs' (i.e. measures of pupil background and school composition). The reported coefficients are multiplied by 100 so as 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 (i.e. Equation (3)); Column (2) shows the same specification estimated on the boundary sub-sample (see Section 3.4) and Column (3) is the cross-boundary (method M2) pair-wise differenced model described in Section 2.3. Columns (4) to (7) introduce the other modifications described in Section 2.3, adding inverse distance weighting (M3), LA boundary dummies (M4), distance-to-boundary polynomial trends (M5), and restricting to boundaries with below-median rates of crossing (M6).

Let us focus first on the price effects of value-added. In the simple OLS estimates, we observe very large and significant associations between school value-added and house prices, with a one point change linked to an 11-15% change in prices (8-11% for a one standard deviation change in the school value-added distribution). These results should not be trusted as causal estimates: in fact, when we eliminate common neighbourhood factors using the boundary differencing strategy there is a dramatic fall in the price effect of school value-added (down to 2%). However, we have argued that the effects of school quality are only separately identified from neighbourhood influences when the distance between matched sales is zero. Therefore, a more reliable estimate is the one presented in Column (4), where we apply IDW weights to the regressions. This shows that the coefficient on value-added rises considerably, up to

3.8%, and becomes more statistically significant. Note that if we follow the strategy of Black (1999) and only concentrate on the closest properties pairs (that is, we apply weights of 1 to transactions within a threshold distance, and weights of 0 otherwise) we find similar results. For example, when we restrict our sample to transaction pairs less than 250metres apart (sample size 16515) we find a point estimate of 3.89, with a standard error of 1.45.

An important result is that once we have applied IDW weights, the coefficient on value-added remains very stable at around 3.7%, (or 3% for one standard deviation) even when we add in boundary dummy variables (Column (5)), and distance-to-boundary polynomials (Column (6)). We can further include boundary × year dummies, instead of simple boundary dummies, to eliminate all time-series variation occurring along boundaries and the coefficients are almost unchanged (3.74 on *va*, 2.75 on *ks1*). Similarly, the results change only slightly when we restrict our analysis to boundaries with low rates of crossing (below median, or less than 5% of pupils crossing along the whole boundary) in Column (7). The size of the house price response sits comfortably with previous results in the literature, surveyed by Gibbons and Machin (2008), which shows a consensus estimate of around 3-4% house price premium for one standard deviation increase in average test scores.

Note that other weighting schemes, for example $e^{-d_{ij}}$ where $d_{ij}$ is the distance between transaction *i* and matched transaction *j*, produce similar results. Additionally, we have experimented with a number of formulations for distance-to-boundary polynomials too, coming to almost identical conclusions. These included: simple difference-in-distance-polynomials (as reported in Table 3); separate polynomials in the distance on the *i* (source) and *j* (matched) sides of the boundary; separate polynomials in the distance of the 'good' and 'bad' sides of the boundary (i.e. an interaction between distance polynomials and an indicator for high or low school value-added). Finally, if we include interactions between distance-to-boundary and boundary dummies, allowing for 680 boundary side specific trends, we find a slightly lower, but still highly significant coefficient on value-added. All in all, our most robust and testing specifications indicate that prices rise by about 3.7-3.8% for a one point increase in school value-added from the mean (about 3% for a one standard deviation change in the school value-added distribution).

Our results also point to a significant relationship between early test scores and housing costs. 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 (again, about 3% for a one standard deviation change in the school age-t test scores distribution). As already mentioned, the interpretation we place on this coefficient is that it measures the house price response due to parental demand for peer group quality. 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 expected peer group quality that are uncorrelated with current academic gains. The net result is that house prices respond to mean age-11 test scores, whether or not these arise through school composition or school value-added. We will return to this point in our Conclusions.

In conclusion, it is worth noting that previous research (Kane et al., 2003 and Gibbons and Machin, 2003) has suggested that single-year test scores could be noisy proxies for the long-run performance indicators in which parents are likely to be interested. This could lead to underestimate the response of prices to expected school performance. In this research, we considered this possibility by using two-year averaged test scores in our regressions, but found no evidence that using single-year performance measures attenuates our coefficients[12].

[12] In fact, the point estimates based on 2-year means come out about 1 standard error lower than the results presented here. Note however that our data set up means that we average over multiple schools in assigning performance to housing transactions. This potentially makes our estimates less sensitive to school specific performance shocks and helps explain why time-averaged measures do not contain more information than single-year indicators.

### 4.4. Falsification tests using imaginary boundaries and inoperative boundaries

In Table 4, we implement the first of our falsification tests based on imaginary boundaries, described as Method 7 in Section 2.3. In the first instance, in Columns (1) to (3), we simply pair sales up with other sales within the same LA, imposing a minimum distance between the matched properties of 20m to achieve better comparability with the actual cross-LA sample. A similar test was carried out in Black (1999). In Column (1), we present the OLS estimates for comparison. In Column (2), we present the coefficients based on the differenced data while in Column (3) we introduce our IDW weighting. Note that we cannot include LA boundary dummies or distance to boundary polynomials in these models, since no boundaries are involved. OLS estimates are similar to what we found before on the full sample. However, when we difference between close-neighbour pairs within the same LA we find no house price effects associated with local schools. This suggests that our findings above are not spuriously driven by local unobservables, rather causally linked to cross-boundary school quality discontinuities.

The specifications based on paired differences across 'fake' LA boundaries drawn 10km North and East tell a similar story. In Column (4), we report simple OLS estimates for comparison. In Column (5), we difference the data across fake LA boundaries, and then go on to apply IDW weights to our regressions (Column (6)) and to include LA boundary dummies and distance-to-boundary trends (Column (7)). The change as we move from Column (4) to (6) is dramatic and illustrates the importance of IDW weighting in our boundary discontinuity design: the simple boundary discontinuity estimates in Column (5) still suggest a significant association of house prices with *ks1* test scores, even when no discontinuity should exist between the school quality assigned to the close-neighbour housing sales pairs (i.e. a similar set of schools could be accessed from both sides of the fake boundary, since these do not act as real barriers). When we apply IDW weights, the coefficients are attenuated and become completely statistically insignificant. In other words, these tests do not falsify our claim that there exists a causal effect on house prices arising from the demand for school quality, when admission is constrained by *real* attendance boundaries. Moreover, the tests provide further support for our use of IDW weighted regressions.

*4.5. Falsification tests using schools which do not admit pupils based on home location*

One way to falsify our findings would be to show that house prices respond to the quality of schools that do *not* ration places according to home address. Our institutional set up allows us to implement this test, as described in Section 2.3 and Section 3.2, using the characteristics of autonomous schools vis-à-vis those of non-autonomous schools. Hence, in Table 5, we compare the effect of school quality on house prices for these two types of institutions (method M8). The first two rows present again the association of house prices with quality in non-autonomous schools, which admit pupils according to home address (i.e. the set of schools used so far for our baseline results). The second two rows show the coefficients for autonomous schools for which home-to-school distance is not an important admission criteria.

In the OLS estimates presented in Columns (1) and (2), we find that the association between school quality and housing prices is large and significant for both types of school, indicating that these coefficients are unlikely to represent causal effects running from school quality to housing demand. In fact, the only reason to buy very close to autonomous schools is to minimise transport costs (not to grant admission). Therefore, the association 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, we find positive and significant results for non-autonomous schools as we did before, but low and insignificant results for autonomous schools – especially when we weight the estimates towards the closest sales pairs (see Columns (3) and (4)). A joint test for the coefficients on value-added and age-7 test scores in Column (4) being equal for autonomous and non-autonomous schools clearly rejects the null hypothesis with a p-value of 0.025.

Once concern is that, given the availability of these two types of schooling, our estimates of the non-autonomous school effects might be attenuated by a tendency for shrewd parents, seeking admission to popular autonomous schools, to buy cheaper housing on those sides of LA boundaries that provide low non-autonomous school quality (and then to ‘cross’ the boundary to attend an autonomous school). Under this scenario, autonomous schools might raise housing prices when non-autonomous quality is low. However, in Column (5) we show that interactions between autonomous and non-autonomous school quality are not significantly linked to prices either, making this hypothesis highly unlikely.

### *4.6. Robustness of the results to sorting, neighbourhood attributes and school resources*

Section 2.4 highlighted the problem associated with inferring mean social valuations of amenities (willingness to pay) such as school quality when households are heterogeneous and there is sorting on school quality according to household type. Figure 4 further showed the fact that some such sorting exists across LA boundaries in our data, in particular for high-qualified residents. Therefore, in Table 6, we check the robustness of our school quality effects to inclusion of a variety of neighbourhood demographic controls (at Output Area level, the smallest geographical unit in the GB 2001 Census containing on average 125 households). We focus in particular on the importance of highly qualified neighbours, with degrees and equivalent qualifications, and those without qualifications. It should be noted these neighbourhood variables are potentially endogenous in these housing price models, because unobserved amenities simultaneously raise housing prices and generate residential sorting.

Column (1) simply repeats our preferred specification from Table 3, while Column (2) adds in a control for the proportion of highly qualified and the proportion of unqualified neighbours. Both enter the regressions with the expected signs and are jointly highly significant, suggesting that households value the educational status of their neighbours (similar to Gibbons, 2003). However, controlling for neighbours’ educational qualifications makes very little difference to the coefficients on school quality. In Column (3) and (4), we go one step further by first adding a range of other demographic controls (Column (3)), and then including the average school achievements of children in the residential neighbourhood (Column (4)).[13] The coefficients on school quality change relatively little, in particular the one capturing the response of house prices to school value-added. This is particularly reassuring since it shows that school effectiveness is capitalised into house prices over and above the educational progress of pupils living in the same neighbourhood. Finally, in Column (5), we subject our data to an even stronger test and match sales across LA boundaries according to whether they are in Census Output Areas in the same quartile of the distribution of high qualifications (in addition to matching on the standard set of housing characteristics). This process provides us with a considerably smaller sample of matched housing pairs, with consequent effects on the precision of our estimates. In fact, the coefficient on *ks1* test scores is weakened considerably, which is consistent with our claim that early test scores act as a proxy for school composition, which is in turn dependent on neighbouring parents' educational background. Nevertheless, our point estimates remain of a similar order of magnitude to our baseline findings, and broadly confirm our results so far. Taken together, the evidence in Table 6 suggests that the second order 'multiplier' effect of school quality on neighbourhood quality operating through residential sorting is quite small and has little bearing on our valuation 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 which 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 M4). It is, however, possible that resources are allocated to LAs in response to changing area

[13] We derive the mean value-added scores of pupils living in the neighbourhood from our pupil database. Neighbourhoods are defined as geographical areas that share the same three nearest schools.

demographics over time, or that localised factors within LAs (e.g. parents' association fund raising) might generate some correlation between within-LA expenditure per pupil and within-LA house prices.

To check the robustness of our findings against these issues, we continue Table 6 by introducing controls for school resources (pupil teacher ratio, expenditure per pupil and pupil numbers) along with a control for local housing tax rates, and by including the school demographic characteristics that affect school income (percentage of pupils eligible for free school meals, ethnic minority proportion and proportion with special educational needs). Clearly, from Column (6), school expenditures, pupil numbers and pupil-teacher ratios show no statistically significant association with prices. This result holds whether or not we control for school value-added or mean test scores, and/or if we replace total expenditure per pupil with sub-categories of spending.[14] More importantly, our key findings on value-added and age-7 test scores are largely unchanged. On the other hand, when we control for other aspects of school composition as in Column (7), the coefficient on age-7 school average test scores falls to near zero and is statistically insignificant. This is mainly because the income-related dimension of intake – namely the proportion of pupils eligible for free school meals – does a better job of measuring those dimensions of school composition that influence parental demand and thus house prices. Other aspects of school composition - ethnicity, special educational needs - turn out to be irrelevant. In contrast, although the coefficient on value-added is attenuated slightly in this rather saturated model, it remains highly statistically significant and economically important in size, emphasising the crucial role of value-added in driving the house price response.

## 5. Concluding discussion

The question of how much parents are willing to pay to get their children into what they perceive as better schools remains a high profile research and policy question. However, accurately pinning down the house price premium generated by superior school performance, and developing a better

[14] This is not surprising given what is known about the weak link between resources and performance that can be observed within cross-sectional data on state school systems - see among others Hanushek (2003) for an international survey and Levacic and Vignoles (2002) for a discussion of the UK experience.

understanding of what aspects of performance parents most value, is hampered by a number of methodological difficulties and concerns.

In this paper, our 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. We started by refining the 'boundary discontinuity' approach to hedonic modelling, and established through a series of novel tests that the methodology provides credible estimates of the causal links between school characteristics and housing prices. These methodological extensions to the boundary discontinuity framework are of broader interest, in that they generalise to other contexts such as border effects in international trade, provision of health care, and the effects of local tax regimes.

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 educational outputs of schools. This is a crucial policy question, because it captures the value of educational performance arising, potentially, from teaching quality, leadership quality and resources. The alternative explanation we considered is that prices rise in response to compositional aspects of schools, which are less amenable to policy intervention and may have little or no bearing on educational effectiveness.[15]

Our results are the first to show convincingly that households pay higher house prices for schools that are likely to *raise* their child's educational achievements – i.e. high-value-added schools. In other words, households pay for the *output* of schooling in terms of expected educational progress. But households pay an additional premium for a favourable distribution of pupil characteristics in these schools – which we represented by higher mean achievements at age 7. In fact, this 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

[15] Empirical studies are mixed in their findings regarding the effect of better peers on individual educational attainment (e.g. Hanushek et al., 2003; Hoxby, 2002; Gibbons and Telhaj, 2008; Lavy et al., 2008). In fact, even if there were significant benefits to be had from 'better' school-mates, these would be capitalised in house prices via school average value-added. Thus, conditional on school 'effectiveness', a significant response of house prices to school composition seems to indicate that parents value better peers even if these generate no observable impact on their child's achievement.

education. On the other hand, ethnic mix in schools does not appear to have an important bearing on prices and the housing market reveals no preference for higher expenditures, generally and on any specific resources, or preferences for smaller classes and schools.

As it turns out, we are unable to say if households *know* exactly what they are paying for. The magnitudes of the effects of school composition and value-added on house prices are similar to each other, so a one point increase in test scores at age 11 is valued the same, irrespective of whether this is achieved through value-added or peer group composition. One potential explanation is that parents use the headline, end-of-primary test results as an indicator of academic effectiveness, but do not have adequate information to differentiate between school results that arise because of high teaching quality, as captured by a higher value-added, and results that come about because the school is 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. However, a second and more likely explanation is that parents value both academic effectiveness and composition aspects of school quality. Either way, the statistical association between school value-added and house prices seems empirically indestructible, regardless of what we do to control for school composition. 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): 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 achievements at age 7. It is useful to benchmark these effects against expected returns and alternative options. Firstly, it is clear that these price responses represent substantial amounts of money, given that the between-school variance in scores is low relative to the variance in achievements across pupils. The price response for a standard deviation in the *pupil* score distribution (2.7 value-added points) is around 11%, or about £20,500 (£1000 annualised) at the house prices prevalent at the time of our study. 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).[16]

Are these figures credible in terms of the value of investment in a child's education? Our answer to this question is positive. To illustrate this, consider that Machin and McNally (2008) estimate the labour market return to a one percentile increase in age-10 test scores, for a cohort of children raised in the 1970s and 1980s, to be about 0.42%. This implies that a one standard deviation improvement in achievement at this age raises future earnings by 12%. In other words, we find that state primary school quality is valued in the housing market at a very similar rate to its expected return in terms of future earnings.

[16] These figures are derived from Independent Schools Information Service web site and available at: http://www.isc.co.uk/FactsFigures_SchoolFees.htm.

## 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", National Bureau of Economic Research Working Paper 11802.

Berry. S, J. Levinsohn and A. Pakes (1995): "Automobile Prices in Market Equilibrium", *Econometrica*, vol. 63(4), 841-890.

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

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.

Card, D., M. Dooley and A. Payne (2008): "School Competition and Efficiency With Publicly Funded Catholic Schools", National Bureau of Economic Research Working Paper 14176.

Cellini, S., F. Ferreira and J. Rothstein (2008): "The Value of School Facilities: Evidence from a Dynamic Regression Discontinuity Design", Princeton University, mimeo.

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. (2008): "Politics, Markets and Schools: Quasi-Experimental Evidence on the Impact of Autonomy and Competition from a Truly Revolutionary UK Reform", University of Florida, mimeo.

Coldron, J. and P. Boulton (1991): "Happiness' as a Criterion of Parents Choice of School", *Journal of Education Policy*, vol. 6(2), 169-178.

Cushing, B. (1984): "Capitalization of Interjurisdictional Fiscal Differentials: An Alternative Approach", *Journal of Urban Economics*, vol. 15(3), 327-356.

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(4), 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), 1-25.

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

Epple, Dennis and Richard Romano (2000): "Neighborhood schools, choice, and the distribution of educational benefits", National Bureau of Economic Research Working Paper 7850.

Fack, G. and J. Grenet (2008): "Do Better Schools Raise Housing Prices? Evidence from Paris School Zoning", Paris School of Economics, mimeo.

Figlio, D. and M. Lucas (2004): "What's in a Grade? School Report Cards and the Housing Market", *American Economic Review*, vol. 94(3), 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 (2), 197-219.

Gibbons, S. and S. Machin (2005): "Valuing Rail Access Using Transport Innovations", *Journal of Urban Economics*, vol. 57(1), 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(1), 99-119.

Gibbons, S, S. Machin and O. Silva (2008): "Competition, Choice and Pupil Achievement", *Journal of the European Economic Association*, vol. 6 (4), 912-947.

Gibbons, S. and O. Silva (2009): "School Quality, Child Wellbeing and Parents' Satisfaction", Centre for the Economics of Education Discussion Paper 103.

Gibbons, S. and S. Telhaj (2008): "Peers and achievement in England's secondary schools", Spatial Economics Research Centre 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(5), pp. 527-544.

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

Hobbs, G. and A. Vignoles (2007): "Is Free School Meal Status a Valid Proxy for Socio-Economic Status in Schools Research?", Centre for the Economics of Education Discussion Paper 84.

Hoxby, C (2002): "Peer Effects in the Classroom: Learning from Gender and Race Variation", National Bureau of Economic Research Working Paper 7867.

Imbens, G. and T. Lemieux (2008): "Regression Discontinuity Designs: A Guide to Practice", *Journal of Econometrics*, vol. 142(2), 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(4), 1603-1637.

Kane, T, D. Staiger and S. Riegg (2005): "School Quality, Neighbourhoods and Housing Prices: The Impacts of Desegregation", National Bureau of Economic Research 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", Centre for the Economics of Education Discussion Paper 102.

Lavy, V., D. Paserman and A. Schlosser (2008): "Inside the Black-Box of Ability Peer Effects: Evidence from Variation in Low Achievers in the Classroom", National Bureau of Economic Research Working Paper 14415.

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(3), 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. and S. McNally (2008): “The literacy hour”, *Journal of Public Economics*, vol. 92, 1441-1462.

Redding, S. and D. Sturm (2008): “The Costs of Remoteness: Evidence from German Division and Reunification”, *American Economic Review*, vol. 98(5), 1766-97.

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(2), 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

| Method: | (1) OLS all England | (2) OLS boundary | (3) Cross-LA boundary M2 | (4) Cross-LA boundary M3 | (5) Cross-LA boundary M4 | (6) Cross-LA boundary M5 | (7) Low crossing sample M6 |
|---|---|---|---|---|---|---|---|
| Age 11-7 Value-added, (year t – t-4) | **10.64 (0.55) | **14.23 (1.03) | **2.06 (0.52) | **3.82 (0.90) | **3.70 (0.87) | **3.69 (0.87) | **3.49 (1.09) |
| Age7 English, maths (year t-4) | **3.66 (0.45) | 0.53 (1.05) | **3.57 (0.52) | **2.86 (0.85) | **2.75 (0.80) | **2.75 (0.80) | **3.07 (0.91) |
| Inverse property distance weights | No | No | No | Yes | Yes | Yes | Yes |
| Admissions authority boundary fixed effects | No | No | No | No | Yes | Yes | Yes |
| Distance to boundary cubic | No | No | No | No | No | Yes | Yes |
| Observations | 1656001 | 138132 | 138132 | 138132 | 138132 | 138132 | 60394 |

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 landcover, number of schools accessible from transaction site, average distance to accessible schools, distance from transaction site to local authority boundary, year dummies. 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 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. Column (7) sample restricted to boundaries with below-median proportions (<5%) of pupils crossing. 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)). Sample based on transaction pairs for second-hand home sales in years 2003,2004, 2005 and first quarter of 2006, from Land Registry "Pricepaid" postcode dataset. 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 tests: Within-admissions zone and fake boundary difference models of the effect of school quality on house prices (Method M7)

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| | OLS within-LA sample | Within-LA | Within-LA | OLS fake LA sample | Cross fake LA | Cross fake LA | Cross fake LA |
| Age 11-7 Value-added, (year t – t-4) | **14.96 | 0.75 | 0.55 | **16.85 | 1.08 | 0.68 | 0.57 |
| | (0.94) | (0.40) | (0.54) | (1.50) | (0.76) | (1.16) | (1.56) |
| Age7 English, maths (year t-4) | **3.28 | *0.74 | 0.79 | -0.328 | **2.74 | 0.24 | 0.15 |
| | (0.83) | (0.35) | (0.48) | (1.83) | (0.67) | (1.23) | (1.23) |
| Inverse distance weights | No | No | Yes | No | No | Yes | Yes |
| Admissions authority boundary fixed effects | - | - | - | No | No | No | Yes |
| Distance to boundary cubic | No | No | No | No | No | No | Yes |
| Observations | 130500 | 130500 | 130500 | 92054 | 92054 | 92054 | 92054 |

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 landcover, number of schools accessible from transaction site, average distance to accessible schools, distance from transaction site to local authority boundary, year dummies. 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 (3) 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 Columns (5) to (7) 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 (3) and (5) to (7)), or clustered on Census ward (Columns (1) and (4)). Sample based on transaction pairs for second-hand home sales in years 2003,2004, 2005 and first quarter of 2006, from Land Registry "Pricepaid" postcode dataset.

Table 5: OLS and cross-boundary difference models. Falsification checks with autonomous schools (Method M8)

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | OLS | OLS | Cross-LA boundary M2/8 | Cross-LA boundary M5/8 | Cross-LA boundary M5/8 |
| Age 11-7 Value-added, (year t – t-4), non-autonomous schools | **9.40<br>(0.51) | **14.46<br>(1.03) | **2.02<br>(0.52) | **3.68<br>(0.87) | **3.70<br>(0.87) |
| Age7 English, maths (year t-4), non-autonomous schools | **2.30<br>(0.43) | -1.23<br>(1.16) | **3.49<br>(0.52) | **2.72<br>(0.80) | **2.72<br>(0.80) |
| Age 11-7 Value-added in autonomous schools | **9.35<br>(0.45) | **9.89<br>(1.05) | 1.07<br>(0.61) | 0.72<br>(0.80) | 0.74<br>(0.89) |
| Age7 English, maths (year t-4), autonomous schools | **7.02<br>(0.43) | **5.76<br>(0.97 | *1.60<br>(0.62) | 0.70<br>(0.80) | 0.66<br>(0.80) |
| Age 11-7 value-added autonomous x autonomous | - | - | - | - | 1.93<br>(1.15) |
| Age 7 English maths, autonomous x autonomous | - | - | - | - | -0.63<br>(0.83) |
| Inverse distance weights | No | No | No | Yes | Yes |
| Admissions authority boundary fixed effects | No | No | No | Yes | Yes |
| Distance to boundary cubic | No | No | No | Yes | Yes |
| Observations | 1656001 | 138132 | 138132 | 138132 | 138132 |

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 landcover, number of schools accessible from transaction site, average distance to accessible schools, distance from transaction site to local authority boundary, year dummies. 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 (5) 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 (3) to (5)), or clustered on Census ward (Columns (1) and (2)). Sample based on transaction pairs for second-hand home sales in years 2003,2004, 2005 and first quarter of 2006, from Land Registry "Pricepaid" postcode dataset.

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

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| | Cross-LA boundary | Cross-LA boundary | Cross-LA boundary | Cross-LA boundary | Cross-LA boundary | Cross-LA boundary | Cross-LA boundary |
| Age 11-7 Value-added, (year t – t-4) | **3.69<br>(0.87) | **3.42<br>(0.89) | **3.12<br>(0.90) | **3.89<br>(1.03) | *2.68<br>(1.20) | **3.12<br>(0.85) | **2.32<br>(0.90) |
| Age7 English, maths (year t-4) | **2.75<br>(0.80) | **2.05<br>(0.79) | *1.85<br>(0.79) | **2.40<br>(0.81) | 1.37<br>(1.12) | **2.37<br>(0.77) | 0.29<br>(0.87) |
| Neighbourhood qualifications | No | p=0.000 | p=0.000 | p=0.000 | Matched quartile | p=0.000 | p=0.000 |
| Augmented neighbourhood controls | No | No | p=0.000 | p=0.000 | No | 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 | No |
| School expenditure | No | No | No | No | No | p=0.336 | No |
| Local housing (council) tax rate | No | No | No | No | No | Yes | No |
| Pupil characteristics | No | No | No | No | No | No | p=0.050 |
| Standard controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Inverse property distance weights | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Admissions authority boundary effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Distance to boundary cubic | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 138132 | 138132 | 138132 | 109941 | 74819 | 137827 | 137655 |

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 landcover, number of schools accessible from transaction site, average distance to accessible schools, distance from transaction site to local authority boundary, year dummies. 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. Sample based on transaction pairs for second-hand home sales in years 2003,2004, 2005 and first quarter of 2006, from Land Registry "Pricepaid" postcode dataset. Neighbourhood qualifications include proportion high qualified and proportion unqualified. Augmented neighbourhood control set includes proportion black, proportion inactive through illness, proportion unemployed, proportion with dependant children, proportion retired. 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 pf pupil eligible for free school meals, percentage of pupils from ethnic minority and percentage of pupils with special educational needs.

Figure 1: Map of Midlands, Manchester and South Yorkshire illustrating admissions district boundary sample

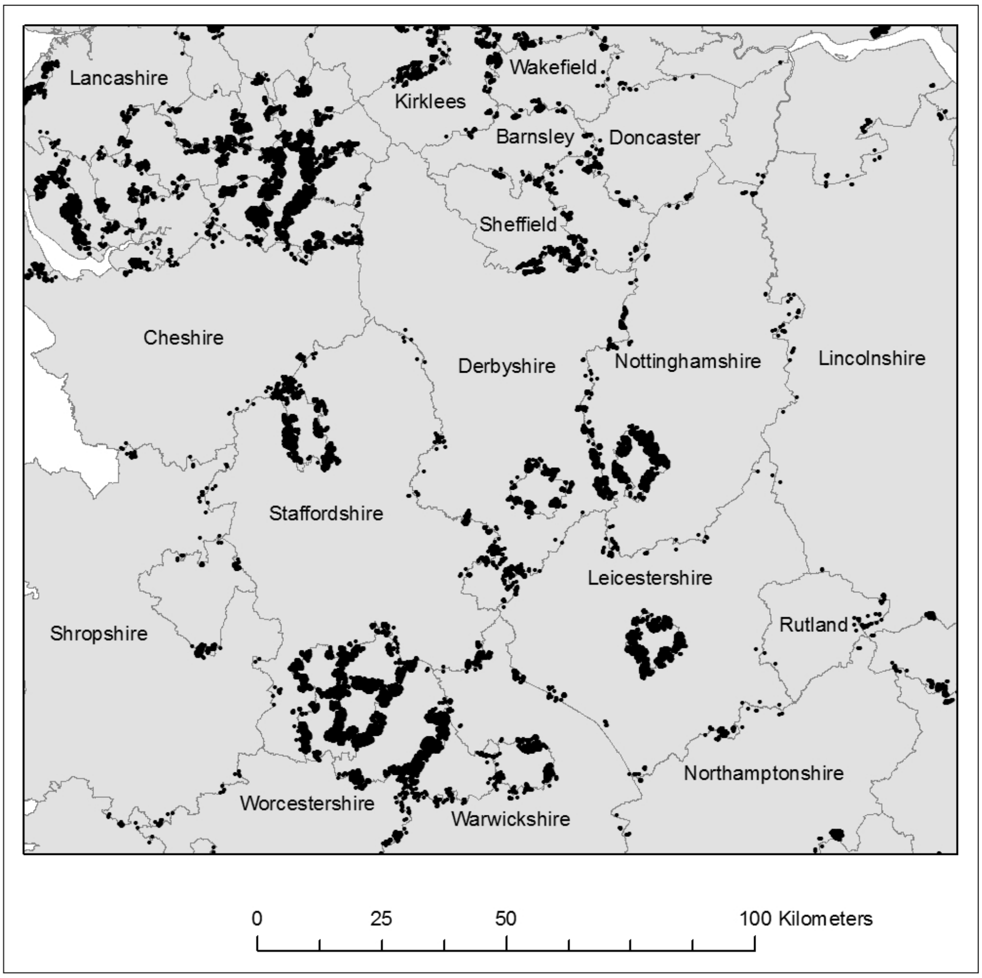

Figure 2: Map of London and the South East area illustrating admissions district boundary sample

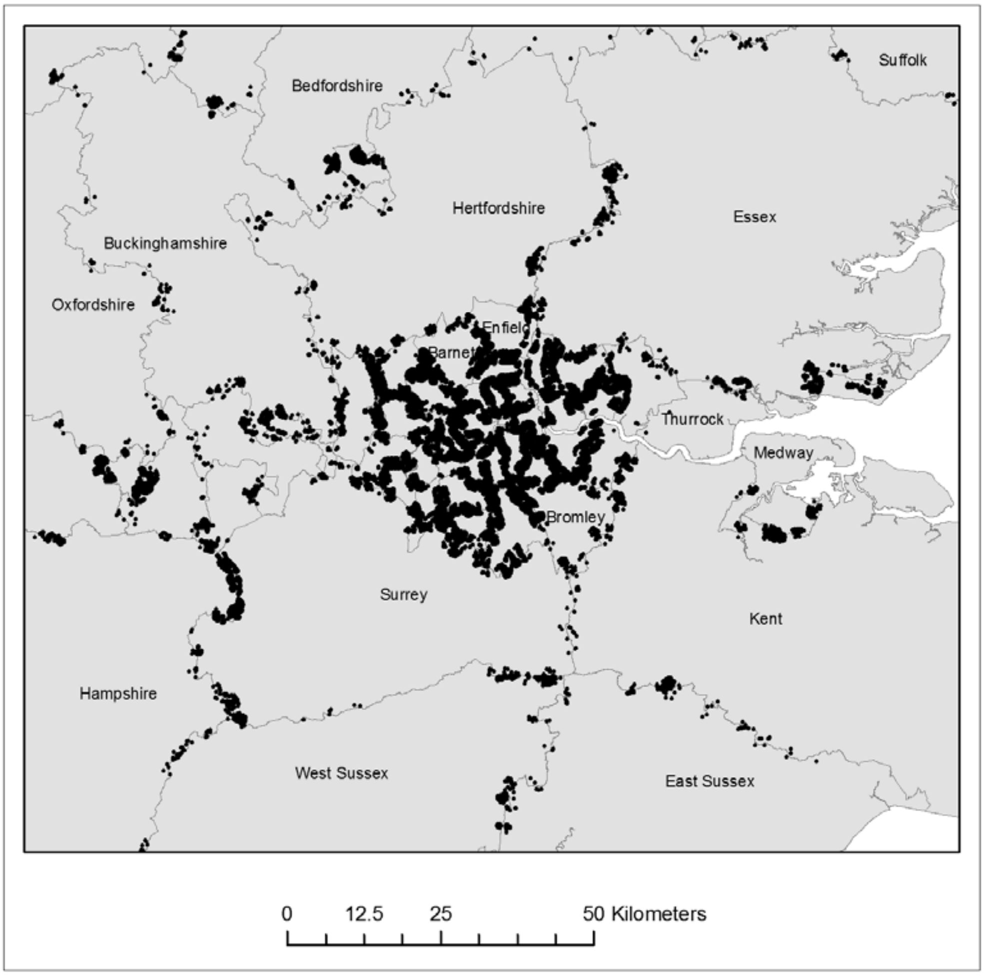

Figure 3: Discontinuities and non-discontinuities in school quality and house prices

*Non-autonomous value-added, by non-autonomous value-added*

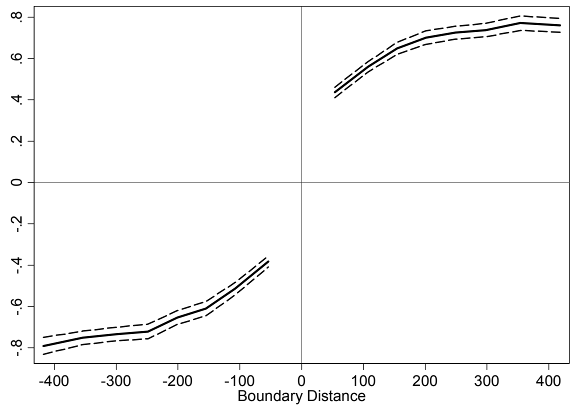


F-tests for equality at boundary:
Two sides equal: 0.000

*Log house price, by non-autonomous value-added*

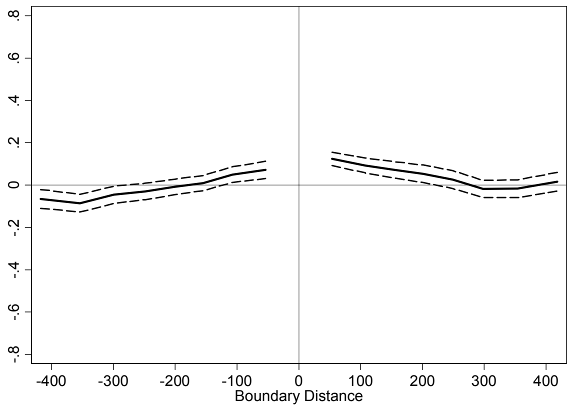


F-test for equality at boundary:
Two sides equal: 0.007

*Autonomous value-added, by autonomous value-added*

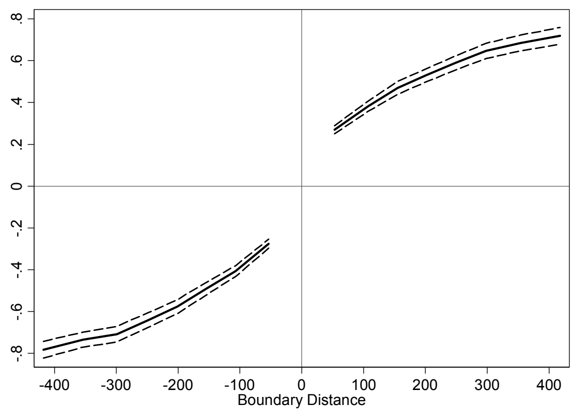


F-tests for equality at boundary:
Two sides equal: 0.000

*Log house price, by autonomous value-added*

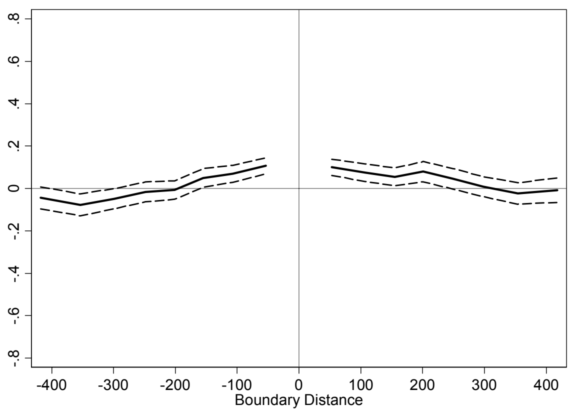


F-tests for equality at boundary:
Two sides equal: 0.760

Figure 4: Discontinuities and non-discontinuities in neighbourhood characteristics

*Households dwelling size*

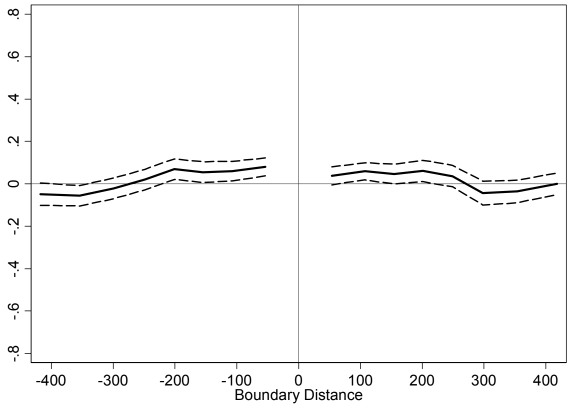


F-tests for equality at boundary:
Two sides equal: 0.326

*Population density*

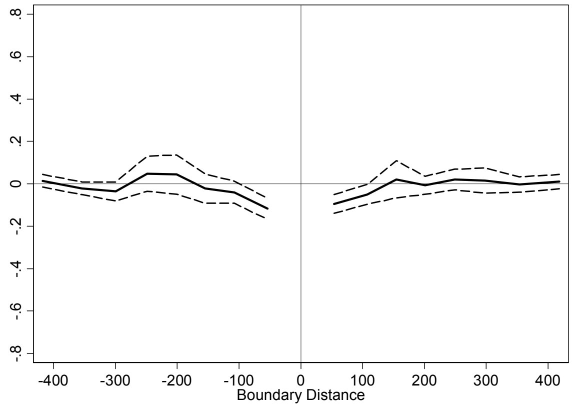


F-tests for equality at boundary:
Two sides equal: 0.946

*Households proportion black*

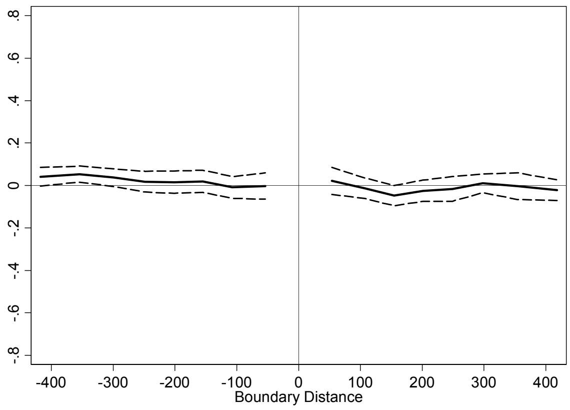


F-tests for equality at boundary:
Two sides equal: 0.914

*Share of dwellings sold*

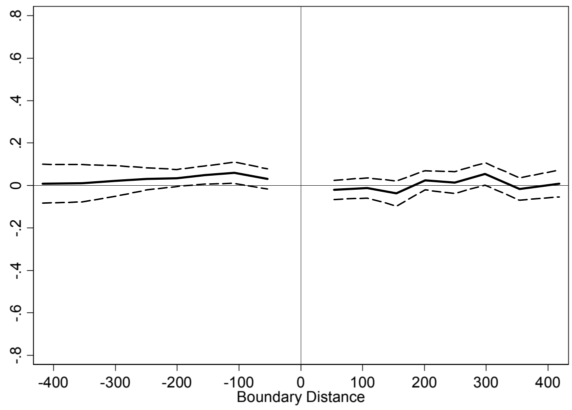


F-tests for equality at boundary:
Two sides equal: 0.076

*Households proportion social tenants*

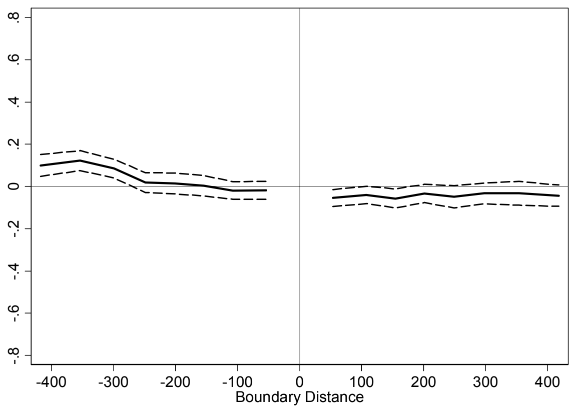


F-tests for equality at boundary:
Two sides equal: 0.116

*Households proportion unemployed*

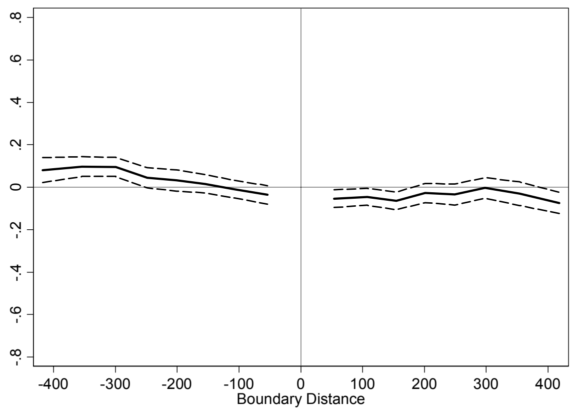


F-tests for equality at boundary:
Two sides equal: 0.142

*Household proportion inactive through illness*

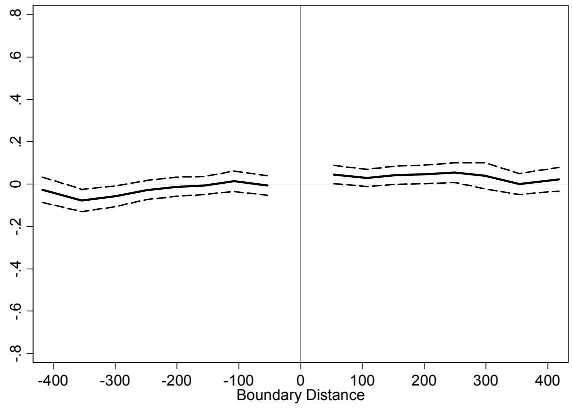


F-tests for equality at boundary:
Two sides equal: 0.105

*Households proportion high qualified*

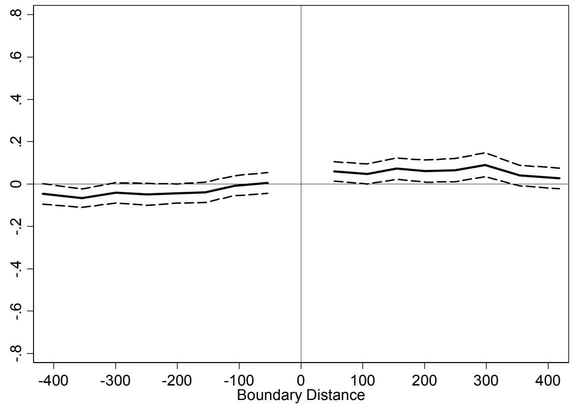


F-tests for equality at boundary:
Two sides equal: 0.007

*Proportion of not retired with dependant children*

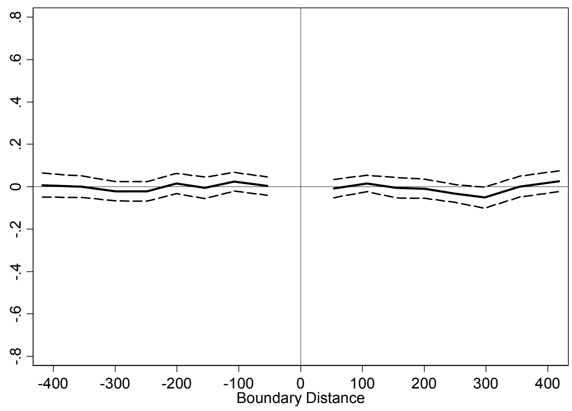


F-tests for equality at boundary:
Two sides equal: 0.660

*Proportion of children in autonomous schools*

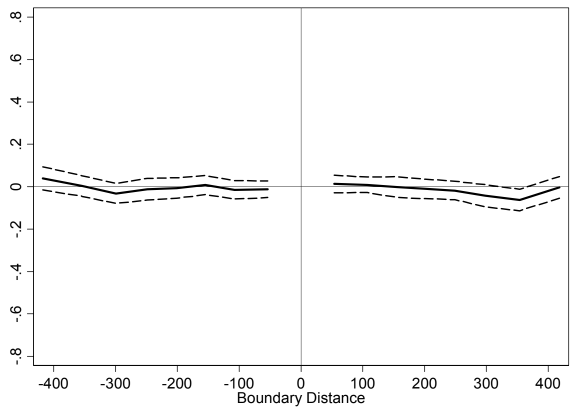


F-tests for equality at boundary:
Two sides equal: 0.290

**Spatial Economics Research Centre (SERC)**
London School of Economics
Houghton Street
London WC2A 2AE

**Tel:** 020 7852 3565
**Fax:** 020 7955 6848
**Web:** www.spatialeconomics.ac.uk

SERC is an independent research centre funded by the Economic and Social Research Council (ESRC), Department for Business, Enterprise and Regulatory Reform (BERR), the Department for Communities and Local Government (CLG) and the Welsh Assembly Government.