The price of free education: Extracting the school quality premium in housing using Brighton and Hove's school admission reforms†

Sajan Shah

Christ's College, University of Cambridge

26th April 2018

Abstract

This paper attempts to extract the school quality premium in housing using a major school admissions reform in Brighton and Hove, the first of its kind in the UK. The 2007 reform abandoned the traditional UK school allocation system based on proximity in favour of a lottery. Using a fixed effects and difference in differences methodology on repeat sales of houses, this paper finds that a 10 percentage point increase in the GCSE pass rate of a school is associated with a 2.38% rise in house prices. This paper is novel as it is the first such paper that applies panel data methods in a quasi-experimental setting to the UK, providing a clean identification strategy that requires fewer assumptions about unobservable household and neighbourhood characteristics than much of the existing literature. It is also unique as the dataset had to be manually constructed using the Geographical Information System software ArcGIS.

Word Count: 7487

†I would like to thank my supervisor, Dr. Julia Shvets, for her valuable feedback, Dr. Gabriel Amable from the Department of Geography for his assistance with my geographical mapping requirements, and Prof. Caroline Hoxby for an insightful discussion.

1. Introduction

There has been much debate in both academic literature and the national media about the relative merits of fee-paying vs. state education, most commonly expressed as a wage-earnings differential. For example, Dolton and Vignoles (2000) find a 7% earnings premium for privately educated individuals within six years of graduating. On the neoclassical assumption that an education investment is worthwhile if the discounted lifetime gain in income outweighs the costs, private school would be the preferred choice for many parents. Despite this, credit-constraints on the poor mean private schools remain reserved for high-income families. This breeds inequality, which from a utilitarian perspective can be damaging to social welfare due to diminishing marginal returns.

Despite the extensive focus on the private vs. state school debate, inequality of access *between* state schools is lesser highlighted because it's not immediately apparent that selection by income could occur in non-fee-paying schools. If access to certain schools is determined by where pupils live, evidence of a school quality premium in housing gives a rational justification to promote an extension of the debate to this sphere. This paper aims to estimate the marginal willingness to pay for an improvement in school quality of non-fee-paying schools through the premium in house prices. Not only is the topic of interest for equality of opportunity arguments, but it is intriguing to be able to place a monetary amount on how much a parent values their child's education. Bogart and Cromwell (2000) have even gone so far as to say that "a community is known by the school it keeps".

Existing work has found there to be a large and statistically significant premium for houses granting access to better quality schools. Early studies were typically based on multivariate OLS techniques controlling for many observable determinants of house prices. However, the existence of unobserved factors (e.g. quality) means that these estimates may suffer from omitted variable bias. More recent work, pioneered by Black (1999), have used boundary discontinuity designs on the assumption that houses either side of a school district boundary share the same unobserved neighbourhood characteristics. Although this goes a long way to control for unobservable factors, if these factors also change discontinuously along catchment area boundaries the problem remains. Additionally, if better quality schools choose to locate in better neighbourhoods there may be reverse causality. Whilst instrumental variable approaches have been used in some studies, their success has been limited by the problem of finding a strong instrument.

This paper uses the 2007 Brighton and Hove school admissions reform as a quasi-experiment in order to estimate the school quality premium in housing. The reform abandoned an existing system which allocated school places according to proximity, in place of a lottery. By using a fixed effects and difference-in-differences methodology with repeat sales of the *same* houses before and after the reform, this paper uses 'within-variation' to better account for unobserved differences in household and neighbourhood characteristics. Additionally, it relies less on arguably strict assumptions as in boundary discontinuity or instrumental variable approaches. The Brighton and Hove reform was the first of its kind in the UK, and this paper is novel as little to no existing empirical work pertaining to the UK have used quasi-experimental methods of identification. This analysis also required the construction of a dataset based on the spatial location of properties, and Geographical Information System's software ArcGIS was used to accomplish this.

This paper finds a significant premium associated with the increased probability of attending a high-quality school. It finds that a 10 percentage-point increase in the proportion of students gaining at least 5 A*'s-C at GCSE is associated with a 2.38% rise in house prices. This is below conventional estimates, although it does suggest selection by income takes place in non-fee-paying schools.

The rest of the paper is organised as follows: Section 2. reviews the existing literature; Section 3. outlines the reform; Section 4. details the dataset's construction; Section 5. describes the estimation strategy; Section 6. discusses the results and extensions; Section 7. employs robustness checks; Section 8. highlights limitations and checks for validity; Section 9. concludes.

2. Literature Review

Empirically estimating the school quality premium is typically based upon Rosen's (1974) hedonic price function. This comes as the result of an optimisation problem where home buyers select both property characteristics and implicit access to local amenities e.g. school quality, subject to their (heterogeneous) preferences and budget constraints. At the optimum, a consumer balances the marginal benefit of improving any one of these factors (e.g. attending a better-quality school) with the additional cost of doing so. In turn, this optimisation problem establishes a "marginal willingness to pay" (MWTP) for each commodity within the composite vector.

To estimate this MWTP, early studies typically used multivariate cross-sectional regressions to control for many observable house and neighbourhood characteristics that may influence house prices, of which school quality is one. For example, Grether and Mieszkowski (1974) use a sample of over 800 sales of family homes in the US in the 1960s and find that moving from a school district in the 50th quality percentile to the 90th percentile raises house prices by 9%.

However, there are two major problems with this estimation strategy. Firstly, there may be many unobserved house and neighbourhood factors (e.g. house quality) that are not accounted for, and which will therefore lead to omitted variable bias if they are correlated with school quality. Secondly there is potential reverse causality, for example if better schools choose to locate in certain areas based on existing neighbourhood attributes e.g. stronger preferences for high quality education. As such, its use has sharply declined in recent work.

The literature evolved in two main ways to tackle the two issues. Firstly, instrumental variables approaches involve finding a variable that is correlated with school quality but otherwise has no impact on house prices. For example, in a study of 350,000 housing transactions in England Rosenthal (2003) uses the random timing of OFSTED school inspections to conclude that a one standard deviation GCSE pass rate increase is associated with a 2% house price premium. However, the lack of success and sparsity of work in this area is indicative of the fact that it has been very difficult to find suitable instruments (Black and Machin 2011). Secondly, the increased availability of precise data has enabled the use of the boundary discontinuity method, most notably pioneered by Black (1999). The key assumption is that either side of a school catchment boundary, neighbourhood characteristics change gradually whilst school quality may change drastically. Therefore, by assuming the two neighbourhoods share the same unobservable characteristics one isolates the effect due to school quality. By removing the omitted variable bias, this method tends to reduce estimates from multivariate cross-sectional studies; using this method, Black (1999) finds that a 5% increase in test scores is associated with a 2.5% rise in house prices. However, this identification strategy rests heavily on its assumptions; if neighbourhood characteristics also change discontinuously across boundaries it still becomes impossible to disentangle the effect due to school quality. Kane, Riegg, and Staiger (2006) show that discontinuous changes in structural and neighbourhood characteristics can coincide with catchment boundaries. In addition, existing work has mainly been applied to US neighbourhoods. If US neighbourhoods are larger in size and/or more readily cross school district boundaries compared the UK, the boundary discontinuity assumption may be less valid in UK housing markets.

Variation in school quality over time means that panel-data methods can be used to difference out the unobserved differences between houses and neighbourhoods. These strategies typically rest on an exogenous change in school quality due to a policy shock or quasi-experiment. By attributing the change in house prices to the change in school quality attended, one is able to identify the associated premium whilst accounting for any unobserved differences between houses. Using repeat sales of houses in the US from 1983-1994, Bogart and Cromwell (2000) find that redistricting of school districts that leads to the loss of a high-quality neighbourhood school leads to a 9.9% reduction in house prices. Machin and Salvanes (2010) exploit a reform that led to the removal of catchment areas in Norway, and find a positive effect of school quality on house prices.

To date there is little to no work using panel-data methods pertaining to the UK, most notably because of the lack of many policy reforms and perhaps "due to the complexity of many schools' admissions arrangements" (Office of the Schools Adjudicator Annual Report 2015). This paper is therefore quite novel as it uses a quasi-experimental approach from a UK policy shock to estimate the premium, and is able to exploit the effects of changes in educational outcomes via repeat sales of houses pre- and post-reform. This means the identification strategy does not rely on strict assumptions about unobservable household characteristics and neighbourhoods. No existing work has considered the effect of the Brighton and Hove reform on house prices; this means the dataset had to be manually constructed using ArcGIS, a Geographical Information System software.

3. Reform

Brighton and Hove is a coastal region in South England with around 250,000 residents. Although, on average, its schools achieve a GCSE pass rate that is broadly in line with the national average (44.5% vs. a national average of 47.6%), the city is suited to this analysis due to the large variance in school quality. This provides the necessary variation in school quality that is required to identify the premium. It is also worth noting that the closeness of means suggests Brighton and Hove bodes well for external validity arguments, discussed in more depth in Section 8.

The pre-reform entry criteria for schools in Brighton and Hove was typical of most UK local authorities. The main method by which school places were allocated was via the home-to-school distance of the primary residence of the applicant (termed a 'proximity oversubscription criterion'). The idea was that local residents should be given preference e.g. to foster a sense of community or reduce travel times. However, the potential for those who have a higher marginal willingness to pay for school quality to move houses may have unintended consequences on who attends the better schools. The reform in Brighton and Hove was specifically designed to be "more equitable and remove the long-standing creation of 'golden halos' of expensive housing that guaranteed school access around popular schools" (Allen et al. 2010). This is because, pre-reform, the clustering of schools in certain areas combined with the allocation of places based on distance meant that some had a guaranteed 'neighbourhood school' while others "were having to travel across the city" (Brighton and Hove Council).

Therefore, "following mounting pressure over fair access to popular secondary schools" (Eastwood and Turvey 2008), the local authority pioneered the lottery to tackle the issue. This was the first such policy in the UK, and it bodes well for identification since it was precisely targeted to mitigate the existence of a premium. Under the new system, catchment area boundaries were drawn up around each school. The proximity oversubscription criterion was abandoned, and instead schools were allocated via lottery; first a priority lottery for all those living within the catchment area, and then a secondary lottery if additional places existed for those living outside. The catchment areas are shown in Figure 1., provided by the Brighton and Hove local authority.

To identify the school quality premium, this paper uses the Dorothy Stringer/Varndean joint catchment area which is shown in the centre of Figure 1. Dorothy Stringer and Varndean are two high-performing schools with GCSE pass rates of 63% and 57% respectively at the time of reform. This is defined as the proportion of students achieving at least 5A*-C grades at GCSE level. The joint catchment area for the new schools was drawn such that many houses to the North of the two schools were allocated Patcham High as their catchment area school, despite being located relatively close in distance to Dorothy Stringer and Varndean. Patcham had a GCSE pass rate of just 28%. This means some houses that would have previously been guaranteed entry into the two high performing schools, are now forced into a low-performing school. Similarly, coastal houses in the South of the catchment area which were previously forced to attend Longhill High (with a pass rate of just 36%), are now located within the joint catchment for the two high-performing schools. Additionally, houses that were previously located on the doorstep of the two high-performing schools and hence guaranteed entry, are now forced to enter a random lottery with all houses in the catchment. Since the two high-performing schools are significantly above the Brighton (and national) average, whilst the two low-performing schools are significantly below par, the discrete change in expected school quality for houses due to the reform provides variation that can be used to identify the associated premium.

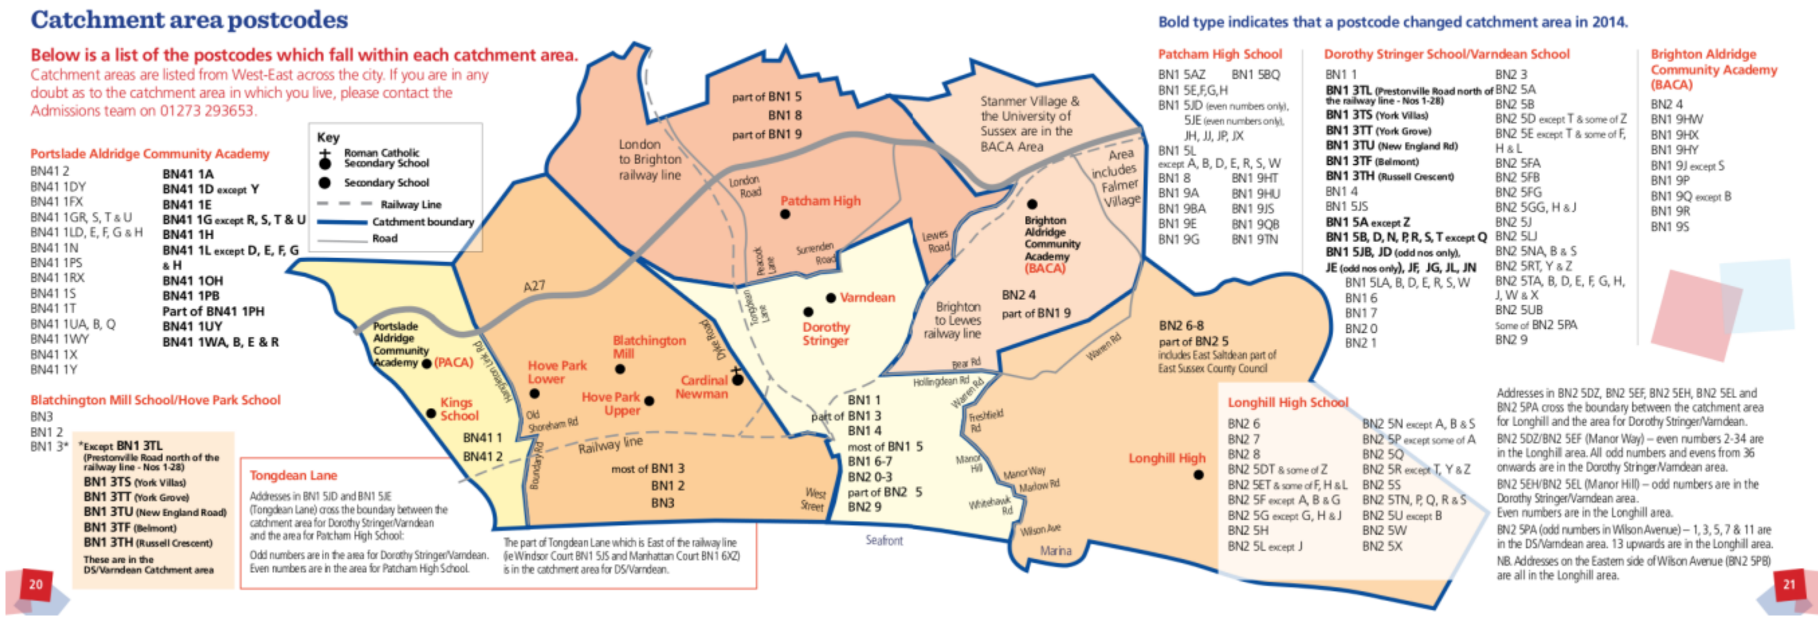

Figure 1: Post-Reform Catchment Areas

4. Constructing the Dataset

School quality can be measured in a variety of ways, such as raw test scores, school expenditure per pupil, or value-added measures that account for differences in pupil intake. Although value-added metrics may be the most accurate measure of the true quality of a school, arguably what matters for the housing premium are parents' *perceived* rather than actual school quality. In the UK, the GCSE pass rate is the most commonly expressed metric e.g. by the press in school 'league tables', and therefore it is reasonable to assume that this is what parents focus on when comparing schools. This is consistent with the work of Downes and Zabel (2002), who find that households judge quality by school outputs (i.e. test results) as opposed to inputs. In addition, positive peer effects that may also be valued by parents are captured using outcomes. This is especially important with regards to Brighton and Hove, since clusters of deprived neighbourhoods means pupil intake varies considerably (Allen et. al 2010). The oversubscribed schools all have the highest GCSE pass rates, so the assumption seems reasonable.

Using house prices to capture the marginal willingness to pay for school quality is justified on two grounds. Firstly, houses are mostly transacted through estate agents, who tend to communicate all relevant information (e.g. on local schools) about the property to the potential buyer. This means that buyers often have complete information, which will be reflected in the price. Secondly, the price of a house represents the expected discounted future utility of anyone who expects to live there, and so even if the house buyer may not derive a first-hand benefit from having access to better quality schools (e.g. if they have no children), they can easily sell the property to someone who does. This means the premium will still exist.

The home sales data covering all residential transactions from 1995-2017 for the Brighton and Hove area was obtained from HM's Land Registry Price Paid Dataset provided by the UK government. This includes 12 years of pre-reform and 10 years of post-reform sales, amounting to over 170,000 transactions across the entire region. The dataset contained the date of each transaction, along with some information on the characteristics of the property being bought/sold. The dataset was initially refined by isolating only houses that were sold *both* pre- and post-reform, and that maintained the same characteristics throughout. For houses with multiple pre- and post-reform sales, only the closest transactions either side of the reform were isolated.

The houses had to be placed onto a geographical map such that expected school quality could be attributed to spatial location. This required the use of the Geographical Information System ArcGIS. The step-by-step process by which the data set was created is outlined as follows, and is illustrated in Figure 2.

Firstly, using Digimap's Ordnance Survey Data, shapefiles of all the postcode boundaries in the area were obtained. Using the postcodes of all the houses contained within the Dorothy Stringer/Varndean catchment area from Figure 1., the relevant shapefiles were isolated and the catchment area was digitally constructed. This is shown in Figure 2i. This enabled the calculation of the post-reform probabilities of getting into Dorothy Stringer or Varndean - for houses within the catchment area this is the lottery probability, and for houses outside the catchment area this is zero since the two schools were oversubscribed.

Since pre-reform allocation was based on proximity, all applicants living within a certain radius R of the schools would be guaranteed entry, and greater than R would be refused entry. This radius was determined as follows. Firstly, the post-reform area of the catchment was calculated. This area was then multiplied by the post-reform lottery probability[1] using data from the Brighton and Hove local authority website. This scaled down the total catchment area to an area within which one would be guaranteed a place i.e. the area of a circle with radius R. R was then simply determined using the area of a circle formula. This is summarised by the following equation:

$$\frac{1}{7}\sum_{r=2008}^{2017}\frac{Number\,of\,Accepted\,Applications_r}{Total\,Number\,of\,Applications_r} \times Post\,Reform\,Area\,of\,Catchment = \pi R^2$$

Inputting the relevant numbers yields R = 1.64 km.

These circles are shown in Figure 2ii.; all houses lying within these circles had pre-reform guaranteed entry to Dorothy Stringer and Varndean, and outside had zero probability.

Finally, the houses were geocoded using the geocoding tool in ArcGIS. Geocoding involves transforming the traditional written address of each property into a set of longitude and latitude coordinates, which are then plotted onto a map. This is shown in Figure 2iii.

[1]The calculation of the lottery probablility is detailed in Table 8. in the appendix.

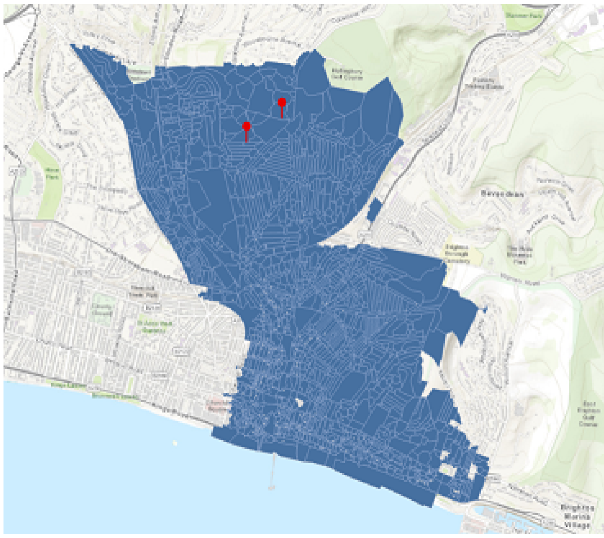

(i) Constructing the Post-Reform Catchment Area

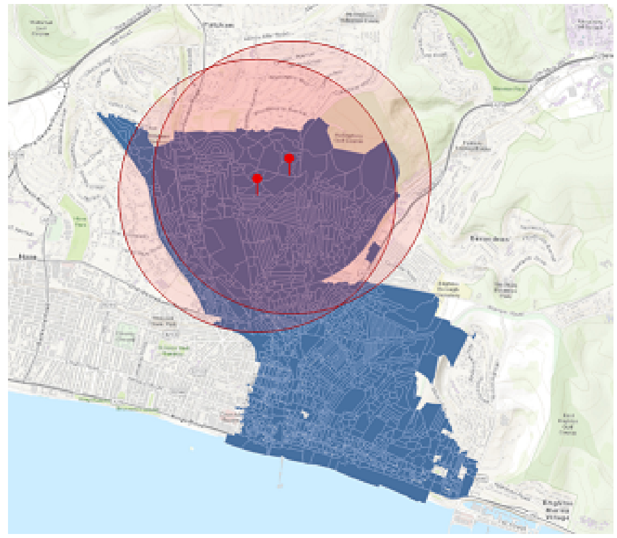

(ii) Calculating the Pre-Reform Radius R

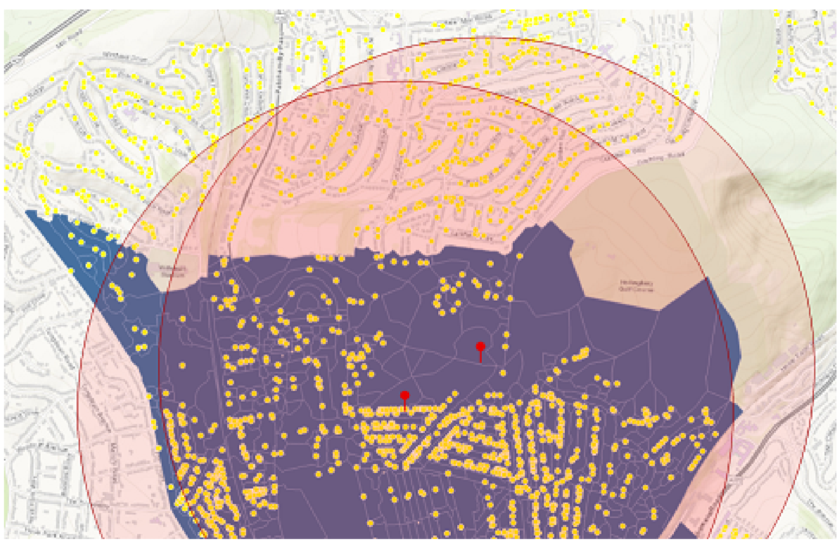

(iii) Geocoding House Addresses

Figure 2: Constructing the Dataset

Based on their spatial location, the houses were split into three treatment groups and a control (detailed in Table 1.) based on how their probabilities of attending the good schools (shown by the red pins) changed due to the reform. Figure 1. provides a geographical illustration.

Group Description	Prob('Good School')[a]		Expected School Quality[b]	
	Pre-Reform	Post-Reform	Pre-Reform	Post-Reform
This is the control group, consisting of houses which lie outside both the post-reform catchment area and the pre-reform radius R, and which are only able to enter Patcham (Control 1) and Longhill (Control 2) both pre- and post-reform. This ensures they have no change in outcomes due to the reform, for example via changes in admission probabilities to other schools not-considered. These houses only ever have access to the under-performing schools, and so have zero probability pre- and post-reform.	0.0	0.0	28% if Patcham High, 36% if Longhill High	28% if Patcham High, 36% if Longhill High
The first treatment group, T1, consists of houses that were previously within the radius R, but are no longer in the catchment area. As a result, they were previously guaranteed entry but now have zero chance. This group mainly consists of the houses to the North of the two schools, which will now likely be allocated Patcham High.	1.0	0.0	61.5%	28%
The second treatment group, T2, consists of houses that were previously within the radius R and are still within the catchment area. This group mainly consists of houses which are relatively close to the two schools. Prior to the reform they had guaranteed entry, but now have only the lottery probability of 0.81.	1.0	0.81[c]	61.5%	55.4%
The third treatment group, T3, consists of houses that were previously outside the radius R, but are now within the catchment area. This group mainly consists of coastal properties to the South of the two schools, and would most likely have attended Longhill pre-reform. These houses previously had no chance of attending the 'good schools', but now have the lottery probability of 0.81.	0	0.81	36%	55.4%

Table 1: Treatment and Control Groups

[a]Prob('Good School') is defined as the probability of attending Dorothy Stringer or Varndean.
[b]See Table 8. in the appendix for a detailed outline of the expected school quality calculations.
[c]This is the lottery probability. See Table 8. in the appendix for an outline of the calculation.

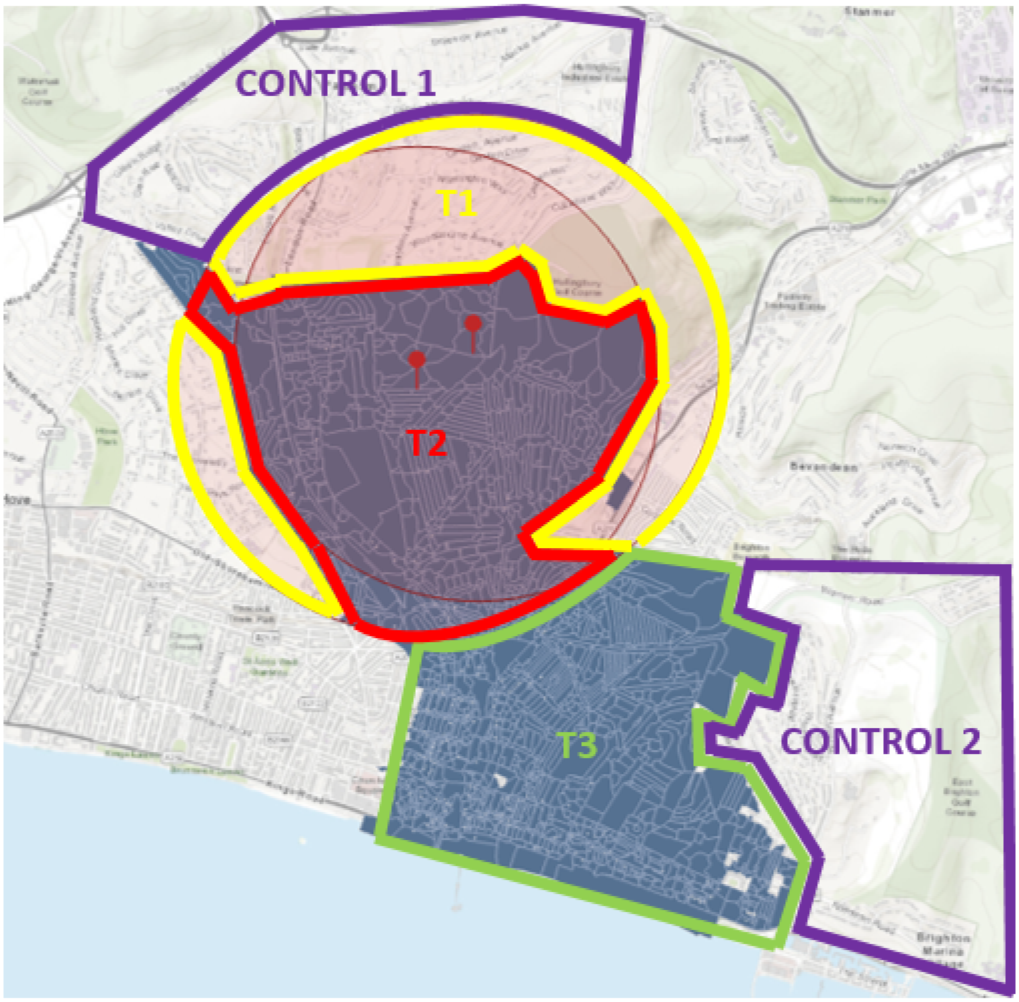

Figure 3: Treatment and Control Groups

Using the "select-by-location" tool in ArcGIS, the houses falling into each area of interest were identified and coded appropriately, creating the final data set. Summary statistics are shown in Table 2.

	Full Sample		**Control**		**T1**		**T2**		**T3**	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Price Paid (£)	244,755	166,893	258,699	130,866	238,819	162,153	241,962	171,912	247,908	173,194
ln(Price Paid)	12.22	0.608	12.34	0.524	12.23	0.553	12.20	0.622	12.22	0.640
Sales Date	01-06-2007		06-05-2007		21-07-2007		10-05-2007		07-07-2007	
House Characteristics (%)										
Detached	7.1		23.6		11.0		4.6		2.4	
Semi-Detached	16.2		54.7		40.5		7.6		2.9	
Terraced	30.1		13.5		11.4		36.6		33.9	
Flat	46.5		8.2		37.0		51.2		60.3	
New Build	1.7		0.6		3.1		1.0		3.2	
Number of Observations	10,412		1,174		1,368		5,554		2,316	

Table 2: Summary Statistics

5. Estimation Strategy

i. Fixed Effects using Repeat Sales

The general hedonic price function (Rosen 1974) for a property i sold at time t is as follows:

$$Price_{it} = B_t \cdot f(S_{it})g(N_{it})h(SchoolQuality_{it-1})k(\varepsilon_{it})$$

where B_t is the regional house price index relative to a base year, S_{it} is a vector of house characteristics, N_{it} is a vector of neighbourhood characteristics and $SchoolQuality_{it-1}$ is the *expected* school quality available to the residents of house i at time *t-1*. School quality is lagged since school preferences at time t will be based on information at the time of application, which is based on results in the previous academic year. For the same house sold at time *t'* (*t'* $>$ t) the equation is identical, except the t subscript is replaced by *t'*.

Using an exponential functional form, we can write this as follows:

$$Price_{it} = B_t \cdot \exp(\gamma S_{it}) \exp(\delta N_{it}) \exp(\beta SchoolQuality_{it-1}) \exp(\varepsilon_{it})$$

Therefore,

$$\frac{Price_{it'}}{Price_{it}} = \frac{B_{t'}}{B_t} \exp\left[\gamma(S_{it'} - S_{it})\right] \exp\left[\delta(N_{it'} - N_{it})\right] \exp\left[\beta(SchoolQuality_{it'-1} - SchoolQuality_{it-1})\right] \exp(\varepsilon_{it'} - \varepsilon_{it})$$

Rewriting in logs,

$$\ln(\frac{Price_{it'}}{Price_{it}}) = b_{t'} - b_t + \gamma(S_{it'} - S_{it}) + \delta(N_{it'} - N_{it}) + \beta(SchoolQuality_{it'-1} - SchoolQuality_{it-1}) + u_{itt'}$$

where $b_t = \ln(B_t)$, $b_{t'} = \ln(B_{t'})$ and $u_{itt'} = \varepsilon_{it'} - \varepsilon_{it}$.

Many existing papers difference away the house and neighbourhood characteristics by assuming they are time-invariant. Since we have a large time horizon from 1995-2017, it is difficult to justify this assumption (e.g. properties may deteriorate over time). However, in order to generate an unbiased estimate of β we only require the assumption that house and neighbourhood characteristics do not change over time in a way that is correlated with the *change* in expected school quality. An in-depth discussion of this assumption, and the checks employed to ensure its validity are detailed in Section 8. By using a panel setup on the *same* houses pre and post-reform, we can account for unobserved differences in house and neighbourhood characteristics even if they are time-varying.

ii. Difference-in-Differences

We also estimate a difference in differences model using the three treatment groups and two control groups from Table 1. This enables us to identify the school quality premium in houses that were affected by the reform in different ways. Via the 'common trends' assumption, any time-varying factors that affect both the treatment and control groups are differenced away. Kuminoff et al. (2010) have shown that the "DID estimator appears to be the best suited to hedonic estimation in panel data".

The following equation is then estimated:

$$\ln(Price_{it}) = \alpha + \beta_1 T_{1_i} + \beta_2 T_{2_i} + \beta_3 T_{3_i} + \gamma Post_t + \lambda_1(T_1 Post_{it}) + \lambda_2(T_2 Post_{it}) + \lambda_3(T_3 Post_{it}) + \varepsilon_{it}$$

where T_n is a dummy variable that takes a value of 1 if the house is in the nth treatment group and 0 otherwise, and $Post_t$ is a dummy variable that takes a value of 1 if the time of sale is post-reform[2] and 0 otherwise.

Based on the nature of the treatment groups, our prior hypothesis is that:

$\lambda_1 < 0$

$\lambda_2 < 0$

$\lambda_3 > 0$

and $\mid \lambda_1 \mid \gg \mid \lambda_3 \mid > \mid \lambda_2 \mid$.

[2]Note this is the time of *announcement*, since house prices will immediately adjust.

6. Results

i. Baseline Results

We first apply the fixed effects and difference-in-differences estimation to the full sample from 1995-2017. Standard errors are clustered at the group level to allow for correlation of errors within each group (Hansen 2007). The results are shown below in column (1) of Table 4. and 5. respectively. Using a continuous quality variable and time dummies, the fixed effects estimation on the whole sample shows that there exists a positive and statistically significant school quality premium in housing. The results show that a 10% increase in the GCSE pass rate of a school is associated with a 1.01% increase in house prices, which is significant at the 1% level. The average house price for the full sample is £244,755, and so on average this amounts to £2448.

The difference-in-differences results show that treatment group T1, which were affected most by the reform, show a large and statistically significant school quality premium; moving from an expected school GCSE pass rate of 61.5% to around 32% is associated with a 10.4% fall in house prices. The results for the other two treatment groups, although consistent with the hypothesis in terms of sign and magnitude, are not statistically significant at the 10% level. To check for consistency between the two results, we compare the coefficients obtained by the difference-in-differences methodology with what would have been predicted by the fixed effects estimation. This is detailed in Table 3. The predicted coefficients based on the fixed effects estimation all lie within the 95% confidence intervals for the difference-in-differences coefficients. This suggests the two methods are consistent with each other.

When examined more closely, Table 3. paints an interesting picture about the two estimations. The fixed effects estimation applied to the full sample predicts a more modest premium for T1 than is generated. Similarly, for T3 it predicts a much larger premium than is generated. This may be justified as follows. If T2 and T3 houses are not exhibiting much of a school quality premium, including them in the full sample fixed effects estimation may dilute the true premium for houses that do show it. There are some potential reasons why this may be the case. For T2 houses, the change in expected school quality is minimal. If the reform does not have a reasonably large effect on education outcomes for these houses, house prices may not respond at all. Additionally, very small changes in house prices due to the small change in expected school quality may be too hard to statistically extract given that the premium is a small in magnitude relative to the value of the house. It may also be argued that buyers of T3 houses may have a hedonic price function that has a low weight on school quality. This means that even if their educational outcomes change substantially, house prices may not respond very much. The reason for this is that T3 houses are coastal properties, and therefore the types of people who are interested in these houses are perhaps not as affected by the choice of schools. For example, if the market for coastal properties is populated by retirees who want to relax by the beach or second-home owners from outside areas, the premium will not feature heavily.

Because of these factors, the fixed effects estimation may be a lower bound since it includes houses for which the premium is not going to show itself very obviously. T1 houses are mainly populated by local families and have a large change in school quality pre- and post-reform, which suggests the premium is likely to be identified. Focusing on this treatment group, the results imply that a 10% rise in the GCSE pass rate is

Group	Change in Expected School Quality	Predicted Difference-in-Differences Coefficient based on Fixed Effects	Actual Coefficient from Difference-in-Differences	95% Confidence Interval for Coefficient	Predicted Coefficient within 95% Confidence Interval?
T1	-33.5%	-0.0338	-0.104	[-0.187, -0.0206]	YES
T2	-6.1%	-0.00612	-0.0102	[-0.0775, 0.0571]	YES
T3	19.4%	0.0196	0.0060	[-0.0691, 0.0811]	YES

Table 3: Checking for Consistency

associated with a 3.10% rise in house prices, or £7,587 for the average house[3]. More simply, going from having access to the two best schools in the area to one of the worst schools is associated with a 10.4% reduction in house prices, or £25,454 for the average house.

ii. Extensions

We decide to extend the baseline estimation in two main ways to tackle potential issues with the initial estimation:

a. Controlling for the Date of Sale

In the difference-in-differences estimation, having a large time period of data both pre- and post-reform may lead to one other difficulty if the houses in the treatment and control groups are being sold at systematically different time periods. Even if on average the sales date were similar in the two groups, if one set of sales dates was a mean-preserving spread of the other there may also be bias. This is because some years may be boom or bust years (e.g. the 2008 Financial Crisis) which may have a significant effect on property values in that year. For example, if a disproportionately large number of houses in a group were sold in 2008 there will be an obvious downward bias. We therefore report results in column (2) and (4) of Table 5. where we normalise all of the sales prices to the year 2015 using the UK house price index. This ensures that any effects due to systematic differences in sales dates between the treatment and control groups are mitigated. This solution also retains the full sample size, which aids with minimising the standard errors.

In the full sample estimation, using the T1 treatment group the results for normalised prices implies that a 10% rise in the GCSE pass rate is associated with a 2.79% rise in house prices, or £8,362 for the average house[4]. Although this is lower as a percentage than the un-normalised estimate, using 2015 as the base year in the normalised sample means that the average house price is higher, so in absolute terms the premium is higher than before.

[3]The average sales price for the full sample is £244,755.
[4]The average sales price at 2015 prices is £299,707.

b. Restricting Sample by Property Type

It is likely that the school quality premium is most likely to exist where potential home buyers will make use of the surrounding schools. The dataset includes the sales of *all* types of residential properties, including flats. It is not unreasonable to assume that flats are unlikely to be populated by families that have children attending secondary schools. The hedonic price function for flat buyers may therefore put a lower weight on school quality, which implies including flats in the sample may lead to attenuation bias. It is worth noting from Table 2. that treatment groups T2 and T3 have a large proportion of flats (51.2% and 60.3% respectively), which may explain their lack of significance so far. To check for this, we estimate the models again only including terraced, semi-detached and detached houses in the sample. These are shown in column (2) of Table 4. and columns (3) and (4) of Table 5.

Contrary to expectations, the estimated premium falls in percentage terms when we restrict the sample to only include terraced, semi-detached and detached houses. Column (4) of Table 5. shows that, using normalised prices, a 10% rise in the GCSE pass rate is associated with a 2.38% rise in house prices. However, since the average sales price has now risen (since flats are no longer included) to £380,566 at 2015 prices, this amounts to £9,057 in monetary terms. This highlights that the absolute premium has *risen*, which is in line with expectations.

VARIABLES	(1) 1995-2017 lnpricepaid	(2) 1995-2017 lnpricepaid
SchoolQuality	0.00101***	0.00071**
	(0.000260)	(0.000313)
Constant	11.03***	11.33***
	(0.163)	(0.149)
Year Fixed Effects	YES	YES
Family Homes Only	NO	YES
Observations	10,412	5,558
Groups	5,206	2,779
R-squared	0.903	0.9275

Robust standard errors in parentheses
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 4: Fixed Effects

VARIABLES	(1) 1995-2017 lnpricepaid	(2) 1995-2017 lnpricepaid	(3) 1995-2017 lnpricepaid	(4) 1995-2017 lnpricepaid
T1	-0.0551*	-0.108***	0.0717***	0.0152
	(0.0304)	(0.0266)	(0.0306)	(0.0248)
T2	-0.127***	-0.160***	0.121***	0.113***
	(0.0268)	(0.0215)	(0.0242)	(0.0196)
T3	-0.0787***	-0.143***	0.177***	0.141***
	(0.0284)	(0.0240)	(0.0295)	(0.0246)
Post	0.602***	0.144***	0.621***	0.145***
	(0.0311)	(0.0276)	(0.0294)	(0.0235)
T1Post	-0.104**	-0.0934**	-0.0872**	-0.0798**
	(0.0431)	(0.0377)	(0.0431)	(0.0353)
T2Post	-0.0102	-0.0238	0.00811	-0.0242
	(0.0337)	(0.0304)	(0.0342)	(0.0278)
T3Post	0.00601	0.00173	0.00252	0.0442
	(0.0392)	(0.0340)	(0.0417)	(0.0348)
Constant	12.04***	12.54***	12.06***	12.59***
	(0.0223)	(0.0195)	(0.0212)	(0.0166)
Normalised Prices	NO	YES	NO	YES
Family Homes Only	NO	NO	YES	YES
Observations	10,412	10,412	5,558	5,558
R-squared	0.225	0.0348	0.329	0.0871

Robust standard errors in parentheses
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 5: Difference-in-Differences

7. Robustness Checks

i. Restricting the Time Period

We decide to restrict the time period of house sales around the policy shock into narrower windows. This serves three main purposes:

i) In a narrower window, it is more likely that the unobserved characteristics are time-invariant, or have not varied in a way that is different between the treatment and control groups. This was the basis behind the 'common trends' assumption. By restricting the time period, this reduces any bias that the failure to satisfy this assumption may introduce.

ii) In order to interpret the 'implicit prices' as the marginal willingness to pay (Kuminoff and Pope 2012), we assume that the ex ante and ex post hedonic price functions are the same i.e. time-constant implicit prices. By restricting the time period, the assumption of time-constant implicit prices is more likely to hold since it is less likely for the hedonic price function to have changed in a narrower window.

iii) Restricting the time period around the reform reduces the influence of any other shocks that may occur at other time periods that may affect the treatment and control groups in different ways.

The original results use sales data from 1995-2017; Table 6. and 7. provide additional results that restrict the data to the periods 2000-2014, 2004-2010 and 2006-2008. However, there is a bias-variance trade-off, since the smaller sample sizes significantly increase the standard errors of the estimates. This is especially a problem since the school quality premium is likely to be a small fraction of the value of the house, and so it becomes harder to disentangle this from random variation. The results show that the coefficients do not change vastly in magnitude as we restrict the time period, although the standard errors increase significantly. Having similar coefficients is reassuring, but the large standard errors render the coefficient in column (2) of Table 6. statistically indifferent from zero at the 5% level, and column (3) at the 10% level.

ii. Varying the Pre-Reform Radius R

The radius R guaranteeing entry into the 'good schools' pre-reform was determined by the method outlined in Section 4. The size of this radius is important because the outcomes for houses either size of the boundary are vastly different - houses within R have guaranteed entry and houses beyond R have no chance at all. It is therefore reasonable to argue that the results may be sensitive to the choice of R. Having said this, matching was a very time consuming process, since the dataset had to be manually constructed. Ideally, we would perform robustness checks on the results with different values of R, however practically this was infeasible due to time constraints.

Despite this, for the T1 treatment group that has been the focus of this analysis, any bias that would arise from using a R that is different from the true R (termed R*) would only serve to reduce our estimates of the school quality premium. To see this, consider the following two cases. If $R < R^*$, the T1 treatment group is smaller than it should be, yet all of the houses contained within it have the correct change in probability. Control 1, however, contains some houses that should have been in T1 (i.e. that experience a fall in expected school quality). But this implies the true effect on T1 is understated since some control group houses would fall in price due to the reform, and offset the measured decrease from T1 houses. If R > R*, Control 1 is smaller than it should be, yet all of the houses contained within it correctly have no change in probability. T1, however, contains some houses that should have been in Control 1 (i.e. experience no change in expected school quality). But this implies the true effect on T1 is also understated, since some of the T1 houses will not drop in price as they have no change in outcomes pre- and post-reform.

VARIABLES	(1) 2000-2014 lnpricepaid	(2) 2004-2010 lnpricepaid	(3) 2006-2008 lnpricepaid
SchoolQuality	0.000946***	0.000863*	0.00139
	(0.000271)	(0.000468)	(0.00147)
Constant	11.57***	12.10***	12.13***
	(0.0178)	(0.0277)	(0.0712)
Year Fixed Effects	YES	YES	YES
Observations	6,972	1,984	354
Groups	3,486	992	177
R-squared	0.817	0.572	0.499

Robust standard errors in parentheses
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 6: Fixed Effects with Restricted Dates

VARIABLES	(2) 2000-2014 lnpricepaid	(3) 2004-2010 lnpricepaid	(4) 2006-2008 lnpricepaid
T1	-0.0892***	-0.137**	-0.131
	(0.0338)	(0.0616)	(0.158)
T2	-0.159***	-0.131**	-0.166
	(0.0264)	(0.0503)	(0.122)
T3	-0.106***	-0.0818	-0.110
	(0.0315)	(0.0545)	(0.135)
Post	0.368***	0.192***	0.187
	(0.0343)	(0.0676)	(0.172)
T1Post	-0.112**	-0.124	-0.0767
	(0.0477)	(0.0875)	(0.220)
T2Post	-0.0262	0.0172	-0.0141
	(0.0389)	(0.0724)	(0.189)
T3Post	0.0133	0.0149	0.0507
	(0.0442)	(0.0781)	(0.197)
Constant	12.19***	12.38***	12.43***
	(0.0262)	(0.0478)	(0.108)
Normalised Prices	NO	NO	NO
Observations	6,972	1,984	354
R-squared	0.139	0.070	0.091

Robust standard errors in parentheses
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 7: Difference-in-Differences with Restricted Dates

8. Limitations and Validity Checks

i. Internal Validity

To give causal interpretation to the measured coefficients, two key assumptions must be made about the treatment and control groups:

1) Common trends between the control and treatment group. In the difference-in-differences analysis, we assume that the treatment and control groups share a common trend to use the control as the counterfactual for the treatment group had they experienced no treatment. This is necessary to ensure that the house and neighbourhood characteristics do not vary over time in a way that is correlated with the change in expected school quality. For example, if the quality of control houses was being improved more quickly than T1 houses, we may falsely attribute the negative coefficient on T1Post to the reduction in expected school quality. Figure 4. checks the common trends assumption by plotting the mean property prices of the control group and treatment groups over time. By visual inspection, there exists no apparent deviation in pre-reform trends which lends confidence that this assumption is likely satisfied.

2) No endogenous changes in group composition. Endogenous changes in group composition in difference-in-difference analyses using a repeated cross-section can often bias results. We overcome this issue by restricting the sample to only include repeat sales of the *same* houses before and after the reform. This means there is no change in group composition at all, endogenous or otherwise.

ii. External Validity

Whether these results can be applied to the UK as a whole depends on how representative Brighton and Hove is. In terms of school quality, Brighton and Hove has an average GCSE pass rate of 44.5%, which is in line with the national average of 47.6%. It also has useful variation in school quality. This suggests that the school decision choice is real and centred around a similar quality level to the rest of the country; it is not the case that parents are choosing between only good or only bad schools. Property prices in Brighton and Hove are second only to London (Office for National Statistics). This may mean that the school quality premium, if taken to be an absolute monetary amount, represents a smaller percentage of the value of houses in Brighton and Hove than the rest of the country. Alternatively, if the school quality premium is a percentage of the property value, the premium in Brighton and Hove is likely higher in absolute terms than the rest of the country.

This paper restricts its attention to *repeat* sales of the same houses pre- and post-reform in the period 1995-2017. It could be the case that houses that are being sold more frequently, and hence more likely to enter the dataset, are fundamentally different from other houses. For example, if families who have a high MWTP for higher quality education, and who currently have children in good quality schools, tend to keep their properties for a long period of time, it may be the case that we are underestimating the true MWTP for the population since these are not included in the dataset. Conversely, it could be that the reform triggered the buying and selling of houses by those who have a high MWTP and who want to move to the best areas, and

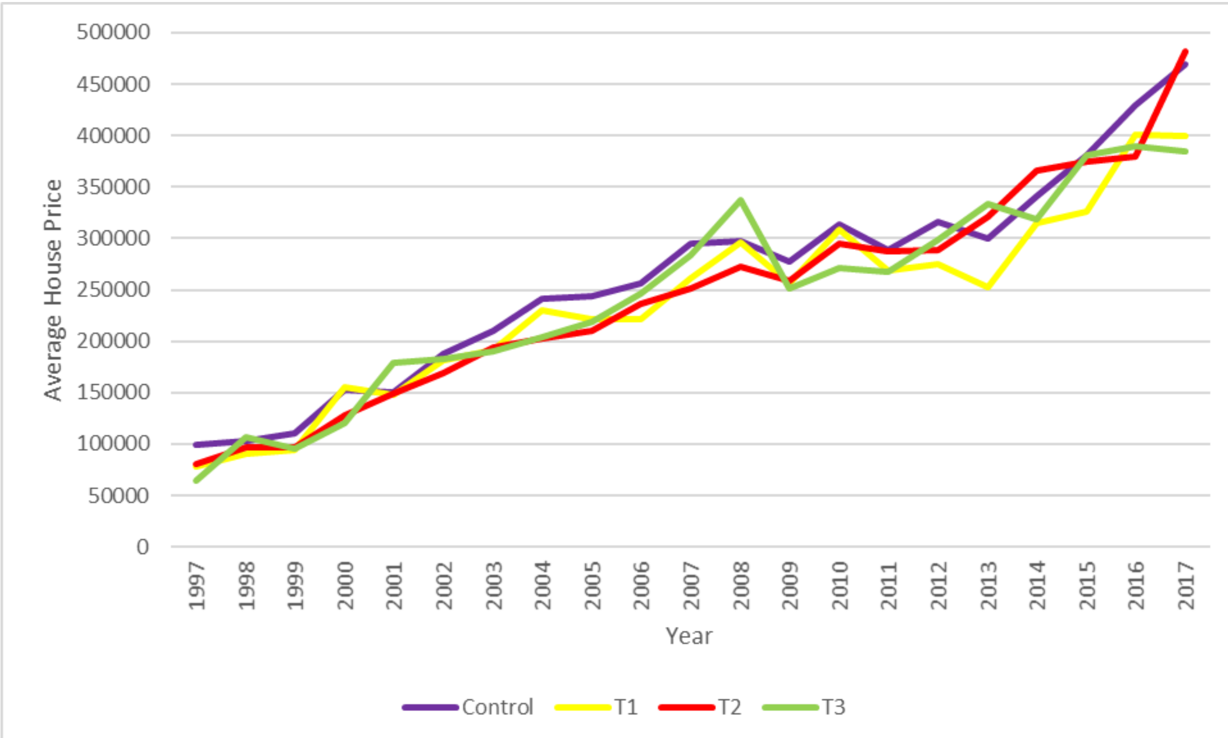

(i) Full-sample trends.

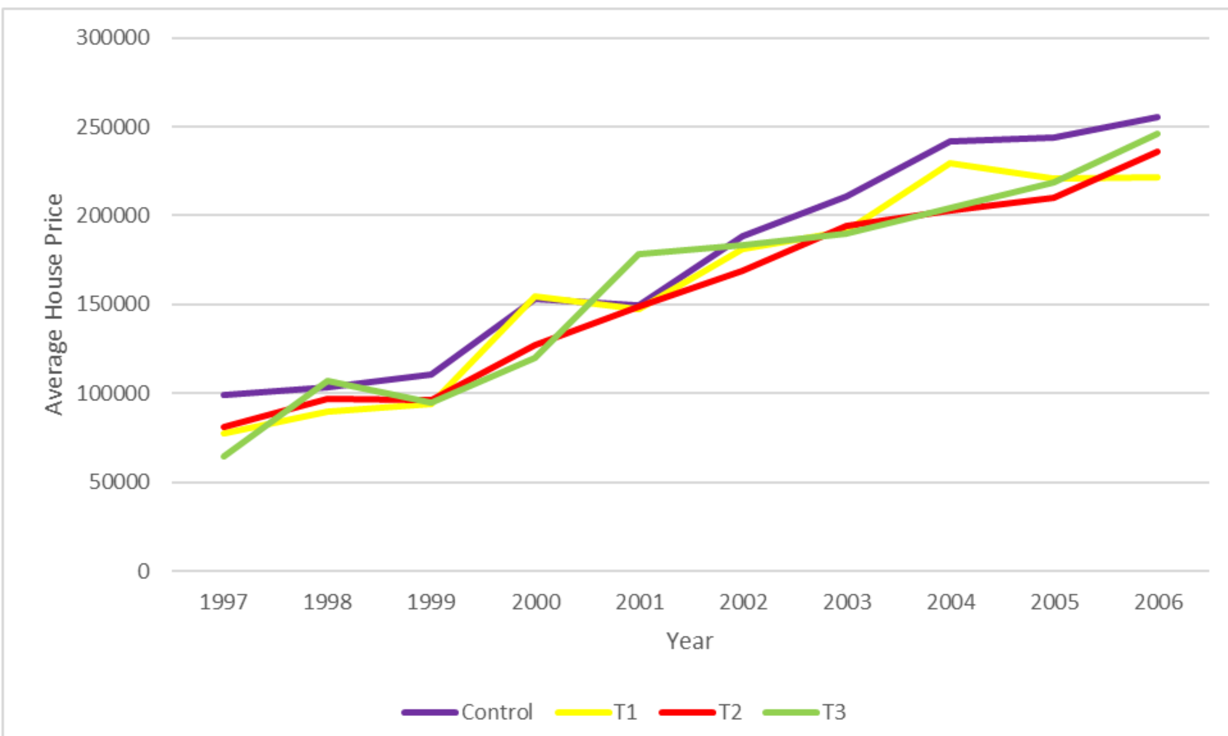

(ii) Graph showing no obvious divergence in pre-reform trends.

Figure 4: Evidence for 'Common Trends' Assumption

so we overestimate the true premium by biasing the sales data towards those who value education highly. Despite these potential issues, it can be argued that the window period is relatively large (over 20 years) and so perhaps any bias is minimal.

9. Conclusions

Using the 2007 Brighton and Hove school admission reform as a quasi-experiment, this paper has shown that there exists a large and statistically significant school quality premium in housing which implies selection by income exists in non-fee-paying schools. It would be useful to consider potential policy implications of this finding in future work[5]. The final difference-in-differences specification finds that a 10% rise in the GCSE pass rate is associated with a 2.38% rise in house prices[6], which amounts to £9,057 in monetary terms when evaluated at 2015 prices. The fixed effects estimation applied to the full sample finds that a 10% rise in the GCSE pass rate is associated with a 1.04% rise in house prices, however this is likely a lower bound. These estimates are slightly below those of most existing empirical work, which may be because this identification strategy is better able to control for unobservable determinants of house prices. Alternatively, this may be because the UK population places a lower weight on school quality in the hedonic house price function compared to the US.

It is important to consider the magnitude of this premium in the wider context of the school choice decision faced by parents. For a single household, the premium amounts to around £9,000 for a 10% rise in the GCSE pass rate. In 2016, the GCSE pass rate for private schools averaged 94.7% (Independent Schools Council 2016), compared to a national average of just 47.6%. Using our results, parents would be willing to pay a £42,000 house price premium for the equivalent difference between two state schools. This is substantially below the fees at private schools, which on average amount to over £112,000 over the course of a seven-year secondary education. In addition, unlike private school fees, the house price premium can be recovered by the homeowner at the point of sale and for families with multiple children, the fixed cost is spread amongst each child. Unless the benefits of a private education extend substantially beyond exam results, these results suggest that by 'paying' to live in areas that guarantee access to the best state schools, parents might be getting a rather sweet deal.

[5] This is constrained by brevity in this paper.

[6] These are the results for the T1 treatment group, which are focused upon.

10. References

i. Literature

[1] Allen, R., Burgess, S. & McKenna, L. (2010) "The early impact of Brighton and Hove's school admission reforms", CMPO Working Paper.

[2] Bailey, M., Muth, R. & Nourse, H. (1963) "A regression method for real estate price index construction", Journal of the American Statistical Association Vol. 58.

[3] Black, S. (1999) "Do better schools matter? Parental valuation of elementary education", The Quarterly Journal of Economics Vol. 114.

[4] Black, S. & Machin, S. (2011) "Housing valuations of school performance", Handbook of the Economics of Education.

[5] Bogart, T. & Cromwell, A. (2000) "How much is a neighbourhood school worth?", Journal of Urban Economics Vol. 47.

[6] Bogart, T. & Cromwell, A. (1997) "How much more is a good school district worth?", National Tax Journal 50.

[7] Brasington, D. (1999) "Which measures of school quality does the housing market measure?", Journal of Real Estate Research 18.

[8] Davidoff, I. & Leigh, A. (2008) "How much do public schools really cost? Estimating the relationship between house prices and school quality", Economic Record 84.

[9] Dee, T. (2000) "The capitalisation of education finance reforms", Journal of Law and Economics.

[10] Dolton, P. & Vignoles, A. (2000) "The incidence and effects of overeducation in the UK graduate labour market", Economics of Education Review Vol. 19.

[11] Downes, T. & Zabel, J. (2002) "The impact of school characteristics on house prices: Chicago 1987-1991", Journal of Urban Economics Vol. 52.

[12] Eastwood, R. & Turvey, K. (2008) "Equal opportunities or loaded dice? The 2007 Admissions Code after the Brighton and Hove adjudication."

[13] Fack, G. & Grenet, J. (2010) "When do better schools raise house prices? Evidence from Paris public and private schools", Journal of Public Economics Vol 94.

[14] Figlio, D. & Lucas, M. (2004) "What's in a grade? School report cards and the housing market", American Economic Review 94.

[15] Gayer, T., Hamilton, J. & Viscusi W. (2002) "The market value of reducing cancer risk: Hedonic housing prices with changing information", Southern Economic Journal Vol. 69.

[16] Gibbons, S., Machin, S. & Silva, O. (2013) "Valuing school quality using boundary discontinuities", Journal of Urban Economics Vol. 75.

[17] Grether, D. & Mieszkowski, P. (1974) "Determinants of real estate values", Journal of Urban Economics Vol. 1.

[18] Hansen, C. (2007) "Asymptotic properties of a robust variance matrix estimator for panel data when T is large", Journal of Econometrics 141.

[19] Kane, T., Riegg, S. & Staiger, D. (2006) "School quality, neighbourhoods and housing prices", American Law and Economics Review 9.

[20] Kuminoff, N., Parmeter, C. & Pope, J. (2010) "Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities?", Journal of Environmental Economics and Management.

[21] Kuminoff, N. & Pope, J. (2012) "Quasi experiments and hedonic property value methods", Working Paper.

[22] Livy, M. (2017) "The effect of local amenities on house price appreciation amid market shocks: The case of school quality", Journal of Housing Economics Vol. 36.

[23] Locke, S. (2013) "Using hedonic and quasi-experimental methods in (dis)amenity valuation with housing data", UKnowledge.

[24] Machin, S. & Salvanes, K. (2010) "Valuing school choice and social interactions: Evidence from an admissions reform", IZA Discussion Paper.

[25] Meese, R. & Wallace, N. (1997) "The construction of residential housing price indices: A comparison of repeat-sales, hedonic-regression, and hybrid approaches", Journal of Real Estate Finance and Economics.

[26] Nagaraja, C., Brown, L. & Wachter, S. (2010) "House Price Index Methodology".

[27] Ries, J. & Somerville, T. (2004) "School quality and residential values: Evidence from Vancouver zoning", Centre for Urban Economics and Real Estate.

[28] Rosen, S. (1974) "Hedonic prices and implicit markets: Product differentation in pure competition", Journal of Political Economy 82.

[29] Rosenthal, L. (2003) "The Value of Secondary School Quality", Oxford Bulletin of Economics and Statistics Vol. 65.

ii. Data

[1] ArcGIS by ESRI, "OS locator: ESRI World Geocoder."

[2] Brighton and Hove local authority, "School admissions information and secondary allocation factsheets." Retrieved from https://new.brighton-hove.gov.uk/schools-and-learning

[3] Brighton and Hove local authority, "Secondary admissions catchment boundaries." Retrieved from https://www.brighton-hove.gov.uk/sites/brighton-hove.gov.uk/files

[4] Digimap, "Brighton and Hove postcode boundary shapefiles." Retrieved from https://digimap.edina.ac.uk

[5] Gov.uk, "Compare school and college performance." Retrieved from https://www.gov.uk/school-performance-tables

[6] HM Land Registry, "HM Land Registry Open Data: 'Price Paid'." Retrieved from https://landregistry.data.gov.uk

[7] HM Land Registry, "UK House Price Index." Retrieved from https://landregistry.data.gov.uk

11. Appendix

Calculation	Method of Calculation	Value
Lottery Probability	This is calculated as $\frac{1}{7}\sum_{r=2008}^{2017} \frac{Number\,of\,Accepted\,Applications_r}{Total\,Number\,of\,Applications_r}$ from within the post-reform catchment area.	0.81
Expected School Quality of 'Good Schools'	$\frac{Number\,of\,Places\,at\,Dorothy\,Stringer}{Number\,of\,Places\,at\,Both\,Schools} \times Average\,GCSE\,Pass\,Rate\,at\,Dorothy\,Stringer + \frac{Number\,of\,Places\,at\,Varndean}{Number\,of\,Places\,at\,Both\,Schools} \times Average\,GCSE\,Pass\,Rate\,at\,Varndean$	61.5%
Expected School Quality for Post-Reform Catchment Houses	$0.81 \times Average\,GCSE\,Pass\,Rate\,at\,'Good\,Schools' + 0.19 \times Average\,GCSE\,Pass\,Rate\,at\,'Bad\,Schools'$	55.4%

Table 8: Calculations