The price of free schools*

Andreas Bjerre-Nielsen† & Mikkel Høst Gandil‡

September 2018

Abstract

A central feature of Scandinavian welfare states is the provision of equal access to free primary education. However, if school performance is reflected in property values, economic inequality may diminish equal access. Using highly detailed geographical data for the universe of sales in Denmark in a boundary discontinuity design, we show that property values reflect the socioeconomic composition of student bodies of primary schools in Denmark. Because attendance boundaries change over time we can validate our baseline estimates under less restrictive assumptions and inspect adjustment dynamics. We document that prices begin to adjust immediately and are fully converged within three years. The estimates indicate that our baseline estimates are not inflated by omitted variable bias. Lastly, we calculate that low income households have to forego between 7 and 10 percent of consumption in order to gain access to a socioeconomically strong school. Our findings underline that even when primary school funding is centralized, there are severe obstacles to ensuring equal access to education regardless of parental background.

1 Introduction

Public provision of primary education is one of the main tasks of local government. While Tiebout (1956) showed that local government may allow for optimal provision of a public and localized good such as education, a theoretical literature has shown that externalities in educational production can lead to inequality and economic inefficiency (Benabou, 1993; Durlauf, 1996). A crucial mechanism is property prices. If prices reflect educational quality, low-income households may be unable to buy access to good schools, which in turn create inequality of opportunity.

In this paper, we provide evidence on the effect of school characteristics on real estate prices in Denmark, a welfare state with extensive transfers between local governments.

*We are grateful for the funding received from The Danish Economic Council of the Labour Movement.

†University of Copenhagen, andreas.bjerre-nielsen@econ.ku.dk

‡University of Copenhagen, mga@econ.ku.dk

Contrary to the American system, the transfers imply that variation in local house prices is not determining school funding to the same degree. Differences in school performance may, therefore, operate more through externalities in education than through funding.

We begin by exploiting cross-sectional variation in school characteristics around attendance boundaries, an approach we refer to as boundary discontinuity design. To decrease omitted variable bias from household sorting around boundaries, we make use of extremely detailed individual geographic administrative data to control for the socioeconomic composition of the 200 nearest adults to a given house sale. Once we introduce these controls, we find that the magnitude of the estimates falls considerably. We show that a standard deviation increase in the school average socioeconomic index is associated with a price increase of between 1.4 and 6 percent, with 3.6 as our preferred estimate. The neighborhood characteristics both affect and are affected by school characteristics. The control variables are therefore partly endogenous and therefore bad controls. We, therefore, argue that the low estimates represent a lower bound of the true effect and that the severity of bad controls is under-appreciated in the literature at large.

We investigate the identification strategy underlying the boundary discontinuity design by recasting it as an IV estimator with a single instrument, a binary variable for a sale being on a specific side of a boundary. Because there is only one instrument, we argue that the exclusion restrictions necessary for interpreting partial estimates of multiple measures of schools in the same regression are unduly restrictive. Such estimates might thus be of little use. Empirically, we show this by replacing the socioeconomic index with ethnic composition as a measure of schools. Due to the very high correlation between the two school measures we get almost indistinguishable results. We argue that the two variables may measure the same underlying latent socioeconomic index. Furthermore, if households do not make the distinction between the two measures, it makes little sense for the econometrician to try to estimate partial effects.

Using school GPA averages we find evidence of positive effects on prices. However, this result is neither robust to the inclusion of neighborhood controls nor controls for the socioeconomic composition of schools. This mixed evidence is somewhat inconsistent with the international literature at large. We show that time variation in GPA within schools is much larger than for socioeconomic variables and therefore argue that parents may use socioeconomic factors as a better proxy for performance.[1]

The boundary discontinuity design is based on cross-sectional variation across geographic space. We complement the existing literature by exploiting time variation in the shape of attendance boundaries. These changes provide exogenous variation in school characteristics, which is not subject to sorting bias. In a difference-in-difference framework, we compare households, which end up in the same district after a change. The

[1]This is consistent with Kane et al. (2003) who argues that yearly school test scores have poor predictive power on house prices.

estimates from this approach are in line with the cross-sectional estimates and somewhat higher than our estimates using the neighborhood controls. This supports our conjecture that neighborhood variables cause attenuation bias due to being partly endogenous. We are therefore confident that we identify a lower bound of the effects of school characteristics on house prices.

With the changes in boundaries, we investigate the time profile of adjustment. We find evidence that prices begin to adjust within the same calendar year after the reallocation to a new school. After three years the prices are at the same level as for houses which have been assigned to the same school throughout.

Our results are important for understanding the implications of residential sorting on equality of opportunity as richer households may buy better educational inputs to the detriment of poorer households. By back-of-the-envelope calculations, we find that a household at the tenth percentile in the income distribution must give up between seven and ten percent of disposable income to move from the tenth percentile to the ninetieth percentile of school socioeconomic composition. This is a major consumption loss and therefore likely to create inequality in educational inputs for children, depending on parental resources.

Attendance boundaries may make the importance of geography larger than need be. The literature on matching and school choice has developed other allocation mechanisms which allow for admittance criteria to be less geographically dependent. However, schools are local in nature as children can only travel a limited distance to receive primary education. We, therefore, argue that geography is a binding constraint and that allocation mechanisms may have to be accompanied by housing and zoning policies to ensure equal access to education.

We begin by reviewing the literature and the theoretical reasons why house prices may reflect valuations of school characteristics in section 2. We then proceed to describe the institutional context of allocating student to Danish primary schools in section 3 and describe our data in section 4. The analysis is presented in section 5 and we perform a back-of-the-envelope calculation of the distributional implications of our results in section 6. Section 7 concludes.

2 Methods and literature

The theoretical link between schools and house prices has been studied in detail by Benabou (1993); Epple and Romano (1998); Durlauf (1996) among many others. These studies investigate the link between residential sorting and the financing and provision of educational services as a local (club) good. An important theoretical link goes through externalities in the educational production function. A simple argument goes as follows: If children benefit from exposure to other children with high human capital and if families with high human capital tend to have more financial resources, then rents will reflect these positive externalities. As high-income families are able to pay more for housing, the low-income families will not be able to buy access to these communities and thus are excluded from the beneficial exposure to strong peers.[2] Importantly, this dynamic can exist even if funding is centralized. The link between schools and housing prices is therefore crucial for the possibility of securing equal access to education.

Empirically, the dominant approach to estimate the importance of schools on prices is to conceptualize a home as a composite good, following Rosen (1974). In equilibrium, such a model equates the price of an amenity to the valuation of the marginal buyer. This is the theoretical underpinning of an attempt to model a hedonic price function, which links local amenities to marginal valuations via house prices.[3]

As with almost all empirical economics, the issue of simultaneity and omitted variable bias is a great concern in this literature. The close links between schools, prices, sorting and local political economy create reverse causality, so that house prices exclude low-income groups from buying houses in a school attendance zone, thereby affecting the student body of the local school and possibly the level of funding. As noted by Black and Machin (2011), empirical research is limited in means to control for this type of effects.[4]

Local amenities cause bias insofar as they correlate with school characteristics and are not included in the regression. These amenities could be physical facilities, such as sports and recreational facilities and access to public transport. Early studies mostly sought to expand the number of controls to reduce the bias, see Kain and Quigley (1970) for an example of this approach. However, the general literature has moved towards trying to reduce omitted variable bias using a more reduced form approach where the source of identifying variation is made explicit. Chief among these approaches is the Boundary Discontinuity Design (BDD), which we also employ in this paper. The identifying assumption for this approach to yield unbiased estimates is that unobserved amenities vary continuously while school characteristics are determined by attendance zones, and thus are discontinuous at boundaries. The unobserved amenities thereby cancel out as they are shared across borders.

Black (1999) was the first to use this approach for school districts in Massachusetts. The number of papers later using this approach are legion. Estimates from BDD are typically five times lower than cross-sectional estimates which shows that unobserved heterogeneity is important in hedonic pricing models, as documented by Kane et al.

[2] See Benabou (1994) for a more rigorous version of this argument.

[3] While the Rosen model most directly lead to a structural econometric approach, Black and Machin (2011) note, that the model forms the basic underpinning of most of the literature on the housing market and schools, though sometimes implicitly.

[4] As argued by Black and Machin (2011) this is especially a problem in the U.S. context, where schools are funded via property taxes. In our context, we investigate within municipality variation, whereby funding is not a problem. However, a nuance is that we in essence measure the after-tax valuation of school characteristics.

(2006); Bayer et al. (2007); Gibbons et al. (2013) among others. A central issue is that school characteristics are not the only thing that changes discontinuously at the border. Bayer et al. (2007) show that estimates in a BDD-framework change substantially once additional controls are included, implying that neighbors and house characteristics also tend to vary at the border. Gibbons et al. (2013) approach this problem with a myriad of control techniques, such as weighting and spatial trends.

A range of papers has shown the importance of other factors for the valuation of schools in a BDD framework. Fack and Grenet (2010) match sales across borders and show that access to private schools diminishes the importance of public school characteristics in Paris. A number of studies explore whether the release of public information about test scores and school-value-added affect the capitalization of school quality in house prices, see Kane et al. (2003); Figlio and Lucas (2004); Kane et al. (2006); Fiva and Kirkebøen (2011); Imberman and Lovenheim (2016).

While the BDD approach has been very popular, the issues with sorting around borders have led to other strategies exploiting temporal variation in either school characteristics or assignment to identify the valuation of schools. Changes to school boundaries have been explored with difference-in-difference approach investigating one-time changes to local school attendance boundaries in Shaker Height, Ohio., U.S. (Bogart and Cromwell, 2000) and Vancouver, Canada (Ries and Somerville, 2010). A drawback in both of these studies is that they have no information about neighborhood quality and school composition and use a single redrawing of the boundaries.[5]

Most studies use test scores as a measure of schools. Quantitative results from numerous studies are reported by Black and Machin (2011). The authors conclude that the baseline estimate is that a standard deviation increase in test scores increase house prices by 3 percent.

2.1 The Boundary Discontinuity Design

We now present the Boundary Discontinuity framework in more detail. We begin by presenting a simple model and then recast the estimator as an Instrumental Variable (IV) approach to elucidate the necessary assumptions to provide causal estimates.

We assume that log house prices of dwelling i is a function of a school characteristic, q_s. Suppressing the time-dimension we write:

$$p_{is} = \kappa + \beta q_s + u_{is}, \tag{1}$$

where κ is a constant. Under the assumption that $E[q_s u_{is}] = 0$ we can estimate (1) by

[5]Ries and Somerville (2010) address this problem using a repeated sales price kernel Additionally, Bogart and Cromwell (2000) acknowledge that their sample is small and limited to only high-quality schools.

regressing p_{is} on q_s. However, the moment restriction can be violated for all sorts of reasons. To the extent that the composition of the housing stock and local amenities correlate with the school characteristics, a simple regression according to (1) will yield biased results.

A first approach is to amend (1) with additional controls for housing characteristics and measures of local amenities. However, in order to yield an unbiased estimate of β, these controls must be exhaustive. A central worry is the role of unobserved amenities, which may correlate with school characteristics. However, as Black (1999) noted, if schools vary discontinuously while unobserved amenities do not, then by comparing houses close to each other, but on either side of the border, unobserved amenities cancel out. In other words, if children are allocated to schools via attendance boundaries (SAB), then this creates a discontinuous jump in school characteristics at the border of two SABs. This insight leads to the Boundary Discontinuity Design (BDD).[6] To put this in formal terms, we add a boundary fixed effect to (1) and exclude the constant term:

$$p_{isb} = \beta q_s + \mu_b + u_{isb}, \tag{2}$$

With the fixed effect we exploit only variation within a border region, b. This implies that we control for all characteristics shared among houses on *both* sides of the border, whether they are observed or not. If $E[q_s u_{isb}|\mu_b] = 0$ we can estimate (2) by way of OLS. To discuss the validity of such assumptions, it is useful to reframe (2) as an IV estimator. Define a dummy r_i, which takes the value of 1 if house i is on the high side of the border b. Under the same assumptions leading to unbiased estimates of (2) it must hold that $E[r_i u_{isb}] = 0$. We can therefore calculate the Wald estimator as:

$$\beta^{Wald} = \frac{E[p_i|r_i = 1] - E[p_i|r_i = 0]}{E[q_i|r_i = 1] - E[q_i|r_i = 0]} \tag{3}$$

Observe that the reduced form, i.e. the nominator in (3), is the average difference in prices across borders. For this to be a valid estimate, we need r_i to uncorrelated with other variables, such as (unobserved) neighborhood characteristics. That is, we cannot allow sorting across the border. In other words, the exact position of the border should be as good as random. This is the standard Regression Discontinuity assumption, that the distribution of covariates is continuous at the discontinuity, see Imbens and Lemieux (2008). If this assumption holds, then the nominator of (3) is the average effect of being on the "high" side of a district border. However, this is not a very useful measure in and of itself, as it needs to be rescaled by the first stage to provide an estimate of a marginal effect. The first stage is the denominator in (3). For the Wald-estimator to

[6] The Boundary Discontinuity Design is equivalent to a Regression Discontinuity Design with distance to a border as the running variable.

yield an unbiased estimate, the exclusion restriction needs to be valid. In other words, r_i must only affect p_i through its effect on q_s. The formulation of the model as a Wald estimator highlights the severity of the exclusion restriction in this framework. In (2) schools are measured by a scalar, q_s. As schools exhibit multiple characteristics, we should ideally have an instrument for each characteristic. Nevertheless, we only have one: the dummy for crossing the border. If being on the high side according to one school measure correlates with being on the high side on some other school measure, then omitted variable bias is still an issue.

Though the inclusion of the border fixed effect removes unobserved factors shared at the border, it does not ensure identification of causal *partial* effects. Thus, the effect of different socioeconomic factors may not be causal when included jointly, as is normally done in the literature. In our data, we observe very strong correlations between a socioeconomic index and ethnicity.[7] Based on the outline above, however, we do not believe that the two factors can be separately identified because both may proxy for the same underlying but unobserved socioeconomic factor. We can therefore use either measure as a proxy for this underlying factor, but not both.[8] Without knowing the underlying correlations between measures, this leaves us little confidence in “horse race” type regressions where multiple measures are included to see which factors explain the most.[9]

Bad controls? As mentioned above, a fundamental issue in BDD is sorting across the boundaries. This is especially a problem with school attendance boundaries. It might be that the marginal buyer considers neighbors with high socioeconomic status (SES) a valuable amenity. Thus, failing to include a measure of the neighbor-composition will bias the estimate of schools. Including the neighbor-composition may, however, cause problems as well. If high-SES schools attract high-SES households, then the neighborhood-SES is a function of the school and will therefore be a “bad control”. However, high-SES households will also send their kids to the local school thereby increasing the school SES further. The measures of schools and neighborhoods are therefore completely intertwined. Nevertheless, we may be able to bound the effect in a simple model.[10]

Suppose that a_i is an amenity of house i (or its' vicinity). We do not observe a_i but have a proxy, $\tilde{a}_i$, which is some function of the school characteristic and the true amenity; $\tilde{a}_i = \pi_0 + \pi_1 q_s + \pi_2 a_i$. Now assume that the true model, instead of (2) is given by the

[7] See section 4 and appendix A

[8] To see why we do not identify the partial effects, observe that children both have an ethnicity and a socioeconomic index. In other words, separate variation in these two characteristics requires different compositions of children. However, a socioeconomic group with a high minority-share is fundamentally different from another group with the same socioeconomic composition but a different minority share. Thus, any unobserved differences between the two groups will bias the partial results.

[9] While we doubt the feasibility of estimating regressions with more than one school characteristic at a time, we do provide results in appendix C in order to ensure comparability with results from the hedonic literature. In the main text, we will not interpret these joint measures.

[10] In this example we follow Angrist and Pischke (2008) closely.

following, where we exclude the fixed effects for convenience:

$$p_{isb} = \beta q_s + \delta a_i + u_{isb}. \tag{4}$$

We assume that $\delta > 0$, $\pi_1 > 0$ and $\pi_2 > 0$. Without including the local neighborhood variable, we would estimate $\hat{\beta}^{+} = \beta + \delta\lambda$, where $\lambda = Cov(q_s, a_i)/Var(q_i)$. This estimate is biased upwards if the amenity correlates positively with the school characteristic. The alternative is to include the neighborhood variable:

$$\begin{aligned} p_{isb} &= \kappa + \beta q_s + \delta a_i + u_{isb} \\ &= \kappa + \beta q_s + \delta \left(\frac{1}{\pi_2}\tilde{a}_i - \frac{\pi_0}{\pi_2} - \frac{\pi_1}{\pi_2} q_s \right) + u_{isb} \\ &= \left(\kappa - \delta \frac{\pi_0}{\pi_2} \right) + \left(\beta - \delta \frac{\pi_1}{\pi_2} \right) q_s + \frac{\delta}{\pi_2}\tilde{a}_i + u_{isb}. \end{aligned} \tag{5}$$

If we include the neighborhood variable, we will therefore estimate $\hat{\beta}^{-} = \beta - \delta\frac{\pi_1}{\pi_2}$, which is negatively biased under the parametric assumptions. In other words, in this simple example we can bound the true effect, β, by estimating models with and without hyper-local neighborhood controls.

If we run a regression of (5), we see that the parameter on the proxy for the amenity is $\frac{\delta}{\pi_2}$. In other words, if we do not observe the true amenity but only the proxy, we cannot identify δ unless $\pi_2 = 1$. Our attempt at bounding β should lead us to respect the fundamental uncertainty about the relation between the proxy amenity and the "true" amenity. If we were to interpret on the magnitude of the coefficient on our proxy, we assume knowledge of the true underlying parameter. We caution that the bad control problem causes issues with both the parameters on schools and the proxy for the unobserved amenity. Due to these considerations, we take a less structural approach to interpretation of parameters on neighborhood characteristics than what is mostly done in the hedonic pricing literature.[11]

2.2 Adjustment to changes in boundaries

Attendance boundaries change from time to time. This entails a shock to the school characteristics for some addresses. In the short term, however, other amenities should be approximately constant. If the estimates we find using the Boundary Discontinuity Design are causal, one should expect effects of the SAB change on the house prices to be of the same sign and magnitude. To investigate this, we employ a difference-in-difference approach. We focus on the closing of a price gap and we therefore reverse the time

[11] As an example, Bayer et al. (2007) interpret the change in parameters on neighborhood characteristic before and after the inclusion of school characteristics as a valuation of neighbors over and above what their effects on school characteristics.

dimension compared to the usual difference-in-difference approach. For school s, we look at dwellings transferred *into* the corresponding SAB. Let r_i be an indicator for whether the address is transferred from one SAB to another at some point and let τ_i be the year the dwelling is transferred. We include a fixed effect for all "arrival SAB"-year combinations and run variations on the following regression:

$$p_{ist} = \theta \cdot r_i \times \mathbf{1}(t < \tau_i) + \lambda r_i + \mu_{st} + u_{its} \tag{6}$$

The fixed effect ensures that we are only using differences within a year within the attendance boundary where all dwellings end up. The parameter λ picks up the time constant difference between those dwellings that are transferred and those who are not. The parameter on the interaction, θ, measures the change in differences before and after the change of the attendance boundary. Naturally, a transfer to a new school can entail a fall or an increase in school characteristics depending on the departure schools. To accommodate this, we let r_i take the value of negative one if a dwelling is transferred into a lower measured school.[12]

The specification in (6) measures the change when transferred from a "low" to a "high" school. This is equivalent to the reduced form found in the BDD-framework. To get at the marginal effects we can reformulate (6) by including the school characteristic as a continuous variable:

$$p_{its} = \psi q_{s't} + \lambda r_i + \mu_{st} + u_{its}, \tag{7}$$

where $q_{s't}$ is the school characteristic in year t of the school s' in the same year. Once again, λ picks up time-constant differences between those transferred and those not transferred, whether they be unobserved or not. After the change in boundaries, school characteristics are the same for all dwellings within the SAB of school s - whether transferred into it or not. Because of the fixed effect, the only variation in school characteristics, therefore, comes from changes in SAB across borders *before* the boundary change and ψ provides an estimate of the effect of school characteristics on prices which we can compare to our BDD estimates. Prices may not adjust instantly, and we, therefore, investigate the timing of responses in the analysis.

The specifications above are simplified for exposition. As we estimate the regressions, we elaborate further on how we control for observed covariates and neighborhood characteristics.

[12] This is equivalent to estimating the model for positive and negative shock and then adding up the "flipped" results, such that the treatment from a negative shock is given a negative value.

3 Primary schools and attendance boundaries in Denmark

Danish primary schools are run by municipalities who decide how to prioritize the general level of funding according to their full set of priorities.[13] Schools are free and parent copayment is forbidden by law. This implies that school funding generally does not vary within municipalities. Students are allocated to schools via residential zones, referred to as school attendance boundaries (SAB), districts or catchment areas.[14] Municipalities can change these boundaries as they wish. Anecdotally, they do this due to projected capacity constraint and development of socioeconomic compositions of schools. Administrative authorities in the municipality usually announce changes in boundaries within a year of implementation. If a child lives within a given boundary, she is guaranteed enrollment in the associated school. If a school is not fully subscribed it is possible for children from other districts to be enrolled. Thus, the living within an attendance boundary is a guarantee, but not a determinant for enrollment into a given school. Bjerre-Nielsen and Gandil (2018a) describe the rules in further detail. Municipalities are financed by income taxes and land taxes which are in general set within a tight bound.[15] An extensive system of transfers between municipalities ensure the municipalities with high expenses due to sociodemographic factors are compensated by other municipalities.

4 Data and descriptive statistics

In this section, we briefly review our data and provide descriptive statistics. Much of this section reflects the same choices and restriction used in Bjerre-Nielsen and Gandil (2018a).

House sales Our main dataset comprises sold dwellings from 2008 to 2015 as provided in the dataset EJSA by Statistics Denmark. We observe price, a date of sale and an identifier, EJENDOMSNUMMER. We link this identifier to individual owners through the dataset EJER. We then link the sales to addresses by merging this data to administrative records on individuals living in their own home from 1990 to 2016. To obtain detailed geographical data we link the addresses to a modified version of the Danish Squarenet. This dataset consists of very small polygons (100$\times$100 meters in densely populated areas). We define the location of a dwelling as the centroid of the associated polygon. This provides

[13]Among other things, Danish municipalities are also responsible for child and elderly care, environmental protection, urban planning and execution of active labor market policies.

[14]In Denmark school districts contain only one school. We, therefore, use the three terms interchangeably.

[15] In 2015 the municipal taxes ranged from 22.5 to 27.8 percent of income, with 50 percent of municipalities within 24.9 and 25.8 percent.

detailed geographic location while maintaining a degree of anonymity, see Bjerre-Nielsen and Gandil (2018c) for documentation. We remove addresses for which we do not have a sufficient degree of precision.[16] We exclude farmhouses and sales where building type is not observed.

School attendance boundaries We obtain school districts for each year from the Central Person Registry (CPR) which publishes a file every quarter with school districts provided by the municipalities.[17] We have data dating back to the nineties. However, due to a municipal reform in 2007, we limit our sample to the period 2008-2015. The districts are voluntarily provided by the municipalities and are not subject to any quality checks. We remove a few municipalities where there are obvious errors in the reporting.[18] The districts are provided as lists of addresses. We clean this data and merge the addresses onto a GIS-dataset containing the spatial features of all plots with associated addresses in Denmark (Martrikelkortet) provided by the Danish Geodata Agency. We define the attendance boundaries as the edge of the spacial polygon made out of the union of the plots in the district. We calculate distances as the shortest Euklidean distance to the boundary from the centroid of the residential polygon. In the boundary discontinuity analysis, we include only addresses with a distance of less than 2000 meters to the borders.[19] To avoid unobserved differences in taxes and provision of public services we only include distances to borders within the same municipality.

School characteristics We apply three different school characteristics. The first characteristic is the mean household socioeconomic index (SES) which we define as the first component a Principal Component Analysis using household income rank and dummies for education and employment of the adult members of the households. We rank the component such that the index is uniformly distributed. The index aligns closely with our intuition that high-SES households are employed, educated and have high incomes. For further detail, we refer to the appendix of Bjerre-Nielsen and Gandil (2018a).

For each school in each year, we calculate the average household-SES of the enrolled students. We call this average the school-SES. In the analysis, this will be our main measure of schools. We also calculate the share of Non-Western immigrants and descendants enrolled in the school. We include this measure as it is very prominent in the public debate in Denmark.[20] The last characteristic is the graduating average of ninth graders

[16]Specifically we remove polygons where the share of the area to the convex hull is below 0.4.

[17]As the CPR overwrites previous versions we a grateful for the due diligence of a retired employee in Statistics Denmark, as this has been the only backup known to us.

[18]These municipalities include Vordingborg and Bornholm. We also exclude Gentofte, as it uses fluid attendance boundaries, which likely is internalized into the location decision of households and thus prices.

[19]We restrict the distance restriction further in parts of the analysis.

[20]As we explain below, we observe such a strong correlation in our data between having a high non-

(I.e. the final year of public school). This average is public information and thus freely available to prospective house buyers from a website. We standardize the school grade averages, such that they have mean zero and a unit standard deviation.[21]

Neighborhood characteristics We link individuals to the Squarenet in order to calculate hyper-local neighborhoods; For every polygon, we calculate the mean of socioeconomic characteristics for the 200 nearest adults in each year. These characteristics are the household-SES index and dummies for employment, long-cycle education and non-Western immigrant or descendant.

Descriptive statistics Descriptive statistics for our sample of sales are presented in Table 1a. The first column display results for the total sample. The following two columns split the sample into two, whether the dwelling is on the low or high side of a border, measured by school-SES. It is evident that the sample is not balanced; on the high side, there is a larger share of single-family homes and higher mean square footage. Conversely, houses on the low side are much more likely to be apartments. These differences are not surprising, insofar as high-SES households are more likely to reside in larger houses and at the same time send their children to the local school. In addition, the neighborhood variables show some unbalances.[22]

In the two right-most column, we calculate statistics for those dwellings, which experience a change in school associations sooner or later and the control group, which maintain the same school association throughout. The transferred dwellings seem overall comparable to the control group. However, there is a larger share of apartments and a higher level of education among neighbors. Half of the sales of eventually changed properties occur prior to the change, reflecting that most of the boundary changes occur around 2011. In Appendix Figure A.1 we report cumulative distributions functions for distances to borders.

Western share and a low school-SES that we do not feel confident in disentangling the two characteristics.

[21] We calculate the z-score of GPA for all available data. However, our sample restrictions cause the mean to be positive and the spread to be smaller than one as evidenced in table 1b.

[22] As previously mentioned these neighborhood variables might be thought of as bad controls, and we will discuss the implications of including them as controls in the analysis.

(a) Dwelling level

Variable	Statistic	All BDD	Low side SES	High side SES	SAB changed	SAB control
M^2	mean	123.04	118.57	127.18	124.76	126.98
	median	119	114	124	120	124
	std	43.05	42.06	43.54	45.90	43.31
Distance to border	mean	880.52	874.21	886.38		
	median	823	814.50	831		
	std	534.07	536.44	531.79		
N: Employment, share	mean	0.79	0.78	0.80	0.80	0.79
	median	0.80	0.79	0.81	0.81	0.81
	std	0.08	0.08	0.08	0.08	0.08
N: Long-cycle education, share	mean	0.11	0.10	0.12	0.16	0.10
	median	0.08	0.07	0.09	0.13	0.07
	std	0.10	0.09	0.10	0.12	0.10
N: Non-Western, share	mean	0.05	0.06	0.04	0.06	0.05
	median	0.03	0.04	0.03	0.04	0.03
	std	0.06	0.07	0.06	0.06	0.06
N: SES, mean	mean	0.55	0.54	0.57	0.58	0.55
	median	0.55	0.54	0.57	0.59	0.55
	std	0.08	0.08	0.08	0.09	0.08
Single-family home	share	0.60	0.56	0.64	0.54	0.66
Terraced house		0.18	0.18	0.17	0.17	0.16
Apartment		0.22	0.26	0.18	0.29	0.18
Before change	share				0.49	
Obs.	Count	341697	164344	177353	11161	383928

(b) School level

	Mean	Median	Std.	N, school-years	N, schools	Within/Total variance
School SES	0.48	0.47	0.11	7473	1357	0.03
Non-Western share	0.11	0.06	0.13	7473	1357	0.02
GPA	0.20	0.21	0.67	5714	984	0.38

Table 1: Descriptive statistics

The table presents selected statistics from the total BDD sample. The prefix "N" denotes that it is a hyper-local neighborhood variable. Only observation under 2000 meters from the border is kept four the boundary discontinuity analysis. We do not impose this restriction when investigating changes in SAB. "Before change" is the share of sales of properties, which are eventually transferred to another SAB but which are observed prior to the transfer. In the analysis, the sample is restricted in different ways described in the text.

We present school level descriptive statistics in Table 1b. As we employ schools in multiple years, we also present the share of variance within school relative to the total variance. This gives us a sense of the stability of school characteristics over time. For school SES and non-Western this share is very low, while almost 40 percent of the variance in GPA stems from variation within schools. Insofar as the socioeconomic composition of schools affects or correlate with student performance, families may regard these socioeconomic statistics as better indicators of school quality than the GPA itself. We do not have data on all schools, and the sample size, therefore, drops whenever GPA is used.

In Appendix Figure A.2 we plot the joint distribution of border differences in school SES and the non-Western share. The correlation is close to -0.9. As we repeatedly state in this paper, we believe that we are not able to identify the effect of ethnicity and socioeconomic status separately with a correlation of this magnitude. We thus doubt the validity of insights from partial effect estimates of SES and non-Western share, holding the other constant. We return to this point in the next section.

5 Analysis

This section is split in two. We first present evidence using a boundary discontinuity design in section 5.1 and then proceed to present the results from the changes in attendance boundaries in section 5.2.

5.1 Static border comparison

We begin by estimating the discontinuities at school borders for the three measures of school, one at a time. First, we construct bins 200-meter bins of distance to the boundary. As is conventional in the literature, we define distances from a border as negative if the address belongs to the school, which has the lower measure of the two schools sharing the border. We then run regressions of the following kind:

$$p_{ibt} = \sum_{d=d^-}^{d^+} \lambda^d \mathbf{1}(dist = d) + \mu_{bt} + \mathbf{X}_{it}\eta + \mathbf{Z}_{it}\delta + \varepsilon_{isbt}, \tag{8}$$

where p_{ibt} is the log price of house i at the border b sold in year t. The border-year fixed effect, μ_{bt} insures that we take out any level differences, shared by the two sides of the border.[23] We include a vector of dwelling characteristics, size and building type, in $\mathbf{X}_{it}$. We also include polynomials of the hyper-local neighborhood variables in the vector $\mathbf{Z}_{it}$. As we do not condition the hyper-local neighborhoods to be on either side of the boundary, we implicitly control for spatial trends across the border by including $\mathbf{Z}_{it}$. We

[23]Note, that sales may be in multiple border regions and thus enter as multiple observations which introduce issues of serial correlation. We therefore cluster standard errors at municipal level.

only estimate (8) for border regions where the difference in school characteristics exceed a standard deviation of the total border difference distribution.[24]

In Figure 1 we plot the estimated λ^ds from the model in equation (8) with SES as the measure of schools. The black points are from an estimation without controls and show a clear discontinuity, with a mean difference of 0.16 points. However, much of this price difference is due to differences in house characteristics. When we include housing characteristics, the difference shrinks to seven points. When we include the hyper-local neighborhoods and house characteristics jointly, the discontinuity falls to three points. This is substantially lower than without controls, though still highly significant.

[24] We only make this restriction for plotting the discontinuities in Figures 1 and 2. In Appendix Figure D.1 we investigate whether the difference in school-SES matter for estimating marginal effects.

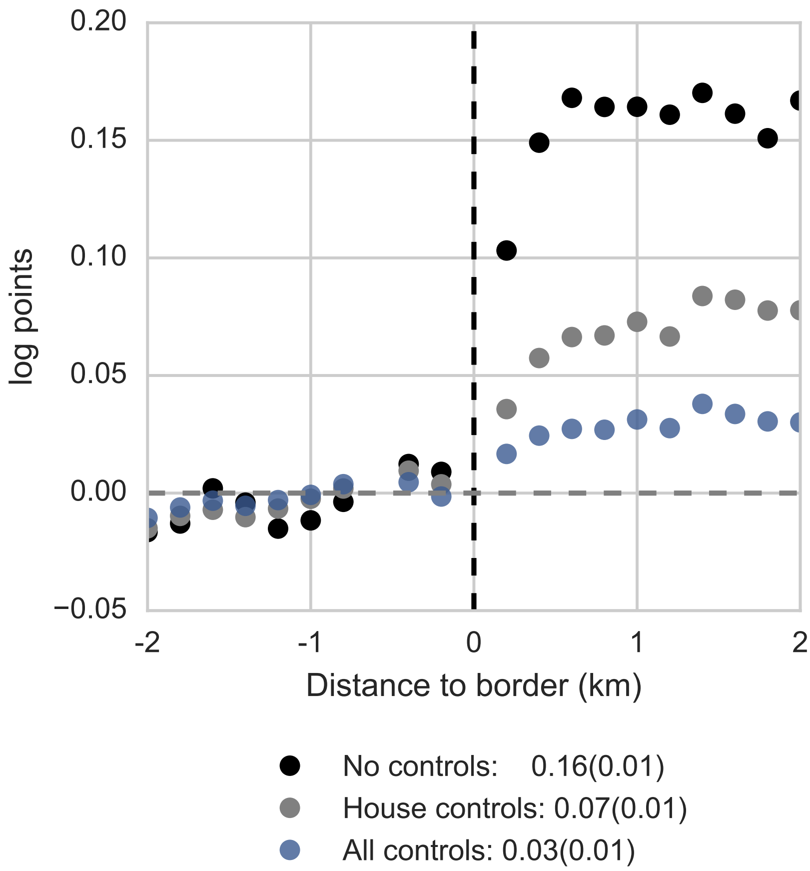

Figure 1: Simple Boundary Discontinuity Design

The graphs show the results of a Boundary Discontinuity Design with SES as school measure. We run regressions of discretized distances to the border, where negative distances signify the address belong to the side of the border with the lowest value of the measure in question. We include a border-year fixed effect to control for level differences shared by both sides of the border. In black, we present the parameters on the binned distance dummies with no additional controls beside the fixed effect. The parameters in gray are from an estimation where we include hyper-local neighborhoods and square meters (including all controls squared). The results are normalized at the 400-meter distance bin at the left side of the border. We also compute an average difference by regressing log prices on a dummy for being on the right side and border-year fixed effects for the same sample. We display the parameters from these regressions in the legend along with standard errors, clustered at the municipal level, in parenthesis. We only include borders with a difference in school SES over one standard deviation. In Appendix Figure B.1 we display corresponding figures with confidence intervals on the binned distance dummies.

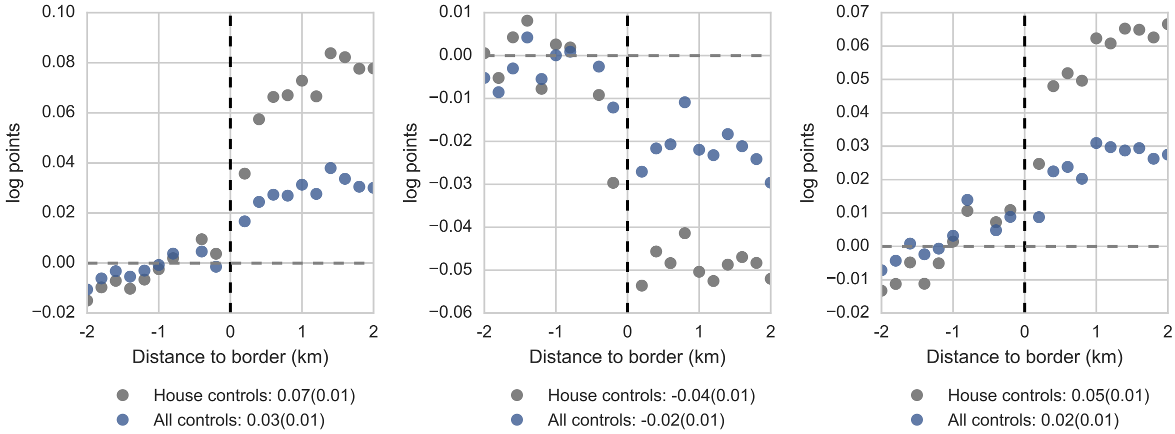

(a) School SES (b) Non-Western share (c) Average GPA

Figure 2: Simple Boundary Discontinuity Design

The graphs show the results of a Boundary Discontinuity Design with one school characteristic at a time. We run regressions of discretized distances to the border, where negative distances signify the address belong to the side of the border with the lowest value of the measure in question. We include a border-year fixed effect to control for level differences shared by both sides of the border. In black, we present the parameters on the binned distance dummies with no additional controls beside the fixed effect. The parameters in gray are from an estimation where we include hyper-local neighborhoods and square meters (including all controls squared). The results are normalized at the 400-meter distance bin at the left side of the border. We only include border regions where the difference is above a tenth of standard deviation of the "border difference distribution". We also compute an average difference by regressing log prices on a dummy for being on the right side and border-year fixed effects for the same sample. We present parameters from these regressions in the legend along with standard errors, clustered at the municipal level, in parenthesis. In Appendix Figure B.1 we display corresponding figures with confidence intervals on the binned distance dummies.

We perform the same exercise for non-Western share and the official GPA-average and exclude the raw estimates without controls as they dwarf the other estimates in magnitude and hinder visual inspection of the discontinuities. Figure 2a repeats the estimation from Figure 1. In Figure 2b we find mimicking results when we measure schools by their non-Western share instead of average SES, though the trends are noisier. The average difference is four log points which decreases to two points when we introduce neighborhood controls. Again, we stress that school-SES and the non-Western share are highly correlated. Lastly, we measure schools by the average GPA and display the results in Figure 2c. We find a discontinuity of five log points in the estimation with house controls. This result falls to two log points, once we introduce neighborhood controls and the visible discontinuity disappears.

As mentioned in section 2, the hyper-local neighborhood may be a bad control insofar as the local sociodemographic makeup is a function of school characteristics. We, therefore, regard these estimates as a lower bound on a true effect. Nevertheless, even if we shut down the effect from neighbors we find significant, positive effects of socioeconomic stronger schools on house prices.

5.1.1 Estimating marginal effects

The discontinuities in Figure 2 provide evidence that schools may causally affect house prices. However, in order to compare these estimates to other results in the literature, we need the discontinuities expressed as marginal effects. We do this by implementing regressions of the following type:

$$p_{ivsbt} = \beta q_{st} + \mu_{vbt} + \mathbf{X}_{it}\lambda + \mathbf{Z}_{it}\delta + \varepsilon_{ivsbt}, \tag{9}$$

where p_{isbt} is the log price of house i in the SAB belonging to s at the border b sold in year t. The school characteristic of school s at time t is measured by q_{st}, and β is the parameter of interest. We include a "house type"-border-year fixed effect, μ_{vbt}. In other words, we are only comparing within housing category within border within year. Due to the fixed effect, all variation in q_{st} comes from crossing the boundary. We restrict our data to be within 300 meters of the district border.[25]

We begin by presenting regressions of prices on one school characteristic at a time where we control for house characteristics, neighborhood characteristics and the border-year-type fixed effects.[26] Column 1-3 in Table 2 present the estimates. We see the same pattern as in figure 2 as school-SES and the non-Western share maintain their signs and significance. A standard deviation of school-SES ($\approx$0.1) causes prices to rise by 1,4 percent. Conversely a standard deviation increase of the NW-share (again $\approx$0.1) cause prices to fall by 0.7 percent. GPA also has a positive effect, but a standard deviation increase in GPA ($\approx$ 0.7) only entails a price increase of 0.3 percent and the effect is very small and insignificant. This effect is an order of magnitude smaller than effects documented elsewhere in the literature.

Column 4 in Table 2 shows the results from regression the three school characteristics jointly. The parameter on a given characteristic is therefore conditional on the other characteristics. The parameter on SES barely changes. The parameter on non-Western share switches sign and falls in magnitude to a no longer significant effect of 0.03. The parameter on GPA remains essentially zero.[27]

[25]In Appendix Figure B.2 we evaluate the importance of the restriction on distance. We see, that while the restriction is important for the raw estimate in column 1 of 2, it hardly matters once we introduce neighborhood controls.

[26]Types include single family homes, terraced housing and apartments.

[27] We note that our estimated null effects of school GPA on local house prices from Table 2 is robust to

	(1)	(2)	(3)	(4)
School SES	0.139***			0.246**
	(0.0345)			(0.0740)
Non-Western share		-0.0733**		0.0561
		(0.0253)		(0.0474)
GPA			0.00432	-0.0119
			(0.00592)	(0.00651)
House controls	X	X	X	X
Neighborhood controls	X	X	X	X
N	56374	56374	52872	52872

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Marginal effect estimates

The models present estimates from an OLS regression of house prices on school characteristics and a border-type-year fixed effects. Standard errors, clustered at municipal level, are presented below the estimates. Only houses within 300 meters of the border enter the regression. For full regression output see Appendix Table B1.

As explained in section 2, we do not believe that our source of variation identifies the partial effects of the three measures. Furthermore, there might be a fourth, unobserved, school characteristic explaining the variation. It is therefore unclear whether it is meaningful to try to distinguish between these variables in the first place. With such high correlations, house-buyers may not make the distinction themselves. The three school measures may essentially reflect the same underlying index from the perspective of the buyers. Thus, trying to separate out the partial effects may be meaningless.[28]

In light of the under-identification of school characteristics, we progress using only the SES as a school characteristic. In appendix table B1 we report the full set of parameter estimates for all school characteristics. In appendix D we investigate heterogeneity in responses as a function of the magnitude of the differences in school characteristics between neighbor schools. We document that differences in prices are almost linearly increasing in differences in school SES, which imply an approximately constant marginal effect of school-SES on prices. For comparison to results in the hedonic literature, particularly Bayer et al. (2007), we construct hedonic regressions in appendix C, but we stress that these regressions are most likely under-identified.

increasing the maximum distances to the boundary, as long as we include school SES and NW, see Table B2 in Appendix B. However, when excluding these school characteristics and increasing the maximum distance to the boundary to 500m the estimate is borderline significant; when including observations within 1000m there is a strongly significant effect. When further excluding neighborhood controls our estimates of GPA on house prices are significant for all maximum boundary distances. In other words, we find a lower bound of essentially zero, but we cannot rule out, that prices might be affected by the average GPA of the local school. The socioeconomic variables, however, seem to carry more weight than the GPA.

[28] Bjerre-Nielsen and Gandil (2018a) discuss this in further detail.

5.1.2 Threats to identification

The large drops in the discontinuities once we introduce controls indicate that the attendance boundaries are not drawn completely at random. We did not expect this to be the case. However, we feel confident that our hyper-local neighborhoods are sufficient to avoid most of the possible omitted variable bias. If houses differ, such that houses on the "high side" of a border are deemed more desirable and thus more expensive, higher income households will also tend to live in them. The hyperlocal neighborhoods, therefore, control for the unobserved features by proxy. For omitted variable bias to still be an issue, the unobserved characteristics must affect the prices in a way that does not influence the socioeconomic composition of buyers. We find it difficult to construct mental models of such sorting. In this context, the local neighborhood controls, therefore, act as proxies for unobserved heterogeneity in house characteristics. As previously mentioned, these finely grained controls may, however, bias the results towards zero, as they are partly determined by the treatment, i.e. the variation in school characteristics. We, therefore, regard the estimates as lower bounds.

Our findings imply that while studies without access to detailed geocoded data on sociodemographic profiles of residents may overstate the importance of schools, schools are important for house prices. We now move to the second identification strategy of this paper, where we exploit time variation in the shape of SABs to validate our results from the boundary discontinuity design.

5.2 Boundary changes and price adjustments

As discussed in Section 2.2 we perform Difference-in-Difference but reverse time. We, therefore, inspect price adjustments to the *new* school association. We begin by regressing sales prices on dummies for time distance to change, covariates and a "arrival-school"-year fixed effect:

$$p_{ist} = \sum_{j=4}^{4} \lambda^j \mathbf{1}(t - \tau_i = j) + \mu_{st} + \mathbf{X}_{it}\eta + \varepsilon_{ist} \tag{10}$$

We are *not* including a dummy for the time constant differences in prices. In other words, if λ^j is zero this means no difference in the price levels between the dwellings already in the arrival district and those arriving. In order to increase power, we recode dummies from one to minus one if the change in school-SES is negative. Figure 3 presents estimates of equation (10). The years before the change in attendance boundaries the coefficients are below zero and significant, while the estimates are close to zero and insignificant after the change. This implies that the price gap between addresses almost close once the addresses become associated with the same school. The adjustment is quick as the

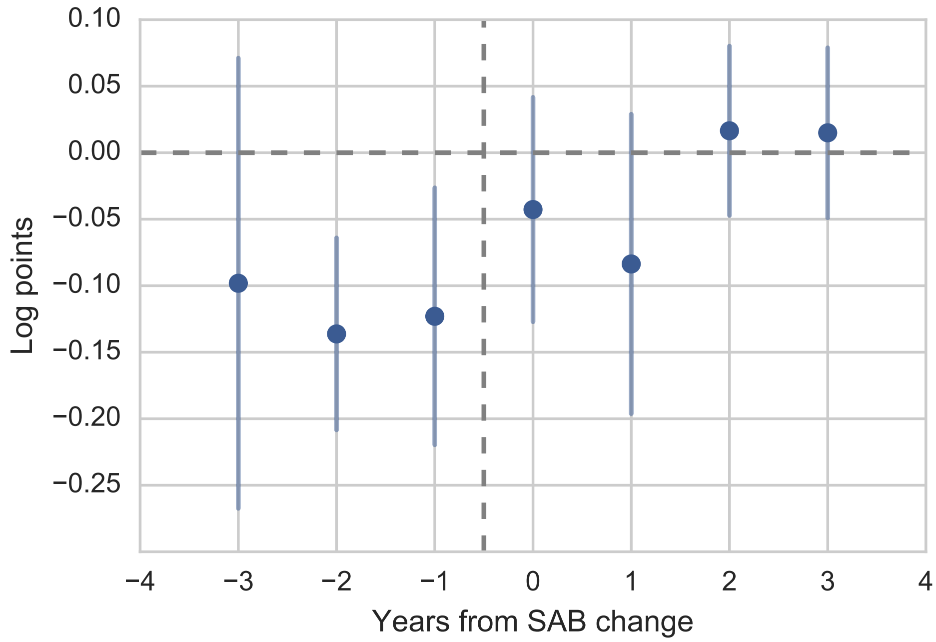

Figure 3: Price adjustment

The figure display estimates of λ^i from Equation (10). Dwelling type interacted with square meters are included as controls. Vertical bars represent 95-percent confidence interval. Standard errors are clustered by arrival SAB.

coefficient rises towards zero already in the year where the boundary change takes place. Three years after the change the prices have completely converged. The adjustment period is somewhat sensitive to the inclusion of controls but the immediate jump in prices suggests that the capitalization occurs quickly. The coefficients provide supporting evidence that the border differences identified in section 5.1 causally affect house prices.

We sum up the differences by collapsing the model to a pre- and a post-dummy. We estimate the model for the time span between four years prior and four years after the change, nine years in total. We estimate the model in equation (6) amended with controls:

$$p_{ist} = \theta \cdot r_i \times \mathbf{1}(t < \tau_i) + \lambda r_i + \mathbf{X}_{it}\eta + \mu_{st} + u_{ist}, \tag{11}$$

where once again μ_{is} is a fixed effect for arrival SAB interacted with year. We continue to define the dummy such that it takes the value of negative one if the 'treatment" is negative, i.e. if the change in boundaries entails a negative change in school-SES. As controls, we include three-way interactions between year of sale, dwelling type and square feet. The estimation results are displayed in panel 3a. Column one displays the results without any controls. Prior to the inclusion into the SAB, the gap was approximately seven log points. Including controls, the effect jumps to ten points in column 2. This may reflect the larger share of apartments in the reassigned group compared to the control group. The estimate remains essentially unchanged when we include neighborhood

(a) Reduced form

	(1)	(2)	(3)
r × Pre	-0.0669	-0.102**	-0.0905**
	(0.0456)	(0.0392)	(0.0332)
r	-0.105*	-0.0195	0.00773
	(0.0408)	(0.0226)	(0.0198)
Cov		X	X
Nbh			X
N	395059	191297	191297

(b) Marginal effect

	(1)	(2)	(3)
School SES	0.778*	0.626*	0.473*
	(0.325)	(0.250)	(0.212)
r	-0.0932*	-0.0254	-0.000363
	(0.0368)	(0.0216)	(0.0189)
Cov		X	X
Nbh			X
N	395059	191297	191297

Table 3: Estimates of reduced form and marginal effects

The top panel display regression results from estimation of (11). The discrete transfers are translated into marginal effects by estimating equation (12). The results of this estimation is displayed in (3b). We do not impose the restriction that households should be within 2000 meters of a border. This explains the higher observation count in column 1 of the two tables. Standard errors in parentheses are clustered by arrival SAB. †$p < .1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

controls as seen in column 3. The parameter on the treatment dummy is close to zero and insignificant once we control for observed dwelling characteristics.

These estimates are well within the range of the BDD-estimates, which are between 3 and 16 log points depending on the inclusion of controls, see Figure 1. Thus, the results provide evidence that the BDD-estimates do not suffer from significant omitted variable bias.

To convert these differences into marginal effects we replace the pre-treatment interaction with the school-SES, in the school associated with dwelling i in period t and run regressions of the following form:

$$p_{ist} = \beta SES_{it} + \lambda r_i + \mathbf{X}_{it}\eta + \mu_{st} + u_{ist}, \tag{12}$$

where λ captures the time constant difference between eventually reassigned dwellings and dwellings which maintain the same association throughout. Due to the fixed effect, the only variation in SES_{it} stems from SAB changes. Depending on controls, the estimates of the marginal effect of school-SES on prices fall between 0.4 and 0.8. These results are somewhat higher than the estimates of the BDD estimation, see Appendix Table B1 for reference. The coefficients on r_i are once again very small and insignificant. This indicates that the effect of school-SES is indeed causal and that the BDD estimates do not suffer from positive bias due to unobserved characteristics of the dwelling or the local neighborhood.

6 Distributional consequences of schooling

To put the estimated effects into context, we perform a simple back-of-the-envelope calculation of the consumption, which households must forego to guaranty enrollment in a high-SES school. We begin by converting house prices to annuities. In appendix B.2 we present evidence that the marginal effect of school-SES on prices (in logs) is approximately constant regardless of the magnitude of border changes. We, therefore, parametrize prices as a log-linear function of school quality, $P = exp(\beta(q - \bar{q}) + p_0)$, where $\bar{q}$ is the average school-SES which we set to 0.5 and p_0 is the reference price of the house. Assuming a T year mortgage with annual payments and a fixed interest rate of i, we can calculate the annuity equivalent of the sales price:

$$a = \frac{i}{1-(1+i)^{-T}} \times e^{\beta\left(q-\frac{1}{2}\right)+p_0}$$

We construct a simple example where a family has the choice between two identical houses (with identical amenities), but within two different attendance boundaries. In other words, p_0 is the same for the two houses. The difference in the implied annuities between an identical house associated with schools s and s' is then given by:

$$a_{s'} - a_s = \frac{i}{1-(1+i)^{-T}} e^{-\frac{\beta}{2}} \left(e^{\beta q_{s'}} - e^{\beta q_s}\right) P_0, \tag{13}$$

where $P_0 = e^{p_0}$. We make the calculation for a single-family home in 2015 prices. Given a average price of 12.000 DKK per square meter, and a 140 square meters in an average single family home, we set P_0 equal to 1,68 million DKK. We assume a loan repayment of 30 years and an interest rate of 4 percent.

We calculate the difference in annuity value from moving from the 10th to the 90th percentile of the school SES distribution, which is a move from 0.36 to 0.62 in school-SES. For β we estimated effects in the range of 0.14 and 0.63.[29] We find that the mean of all estimates is 0.36 which we regard as our preferred estimate. We calculate annuity payments from this estimate of β as well as an upper and a lower bound of the cost difference.

The calculated bounds are displayed in Table 4. We calculate a lower bound of 514 USD per year while the upper bound is 2,308 USD. Our preferred estimate is a yearly

[29] In order to assume that p_0 is the same for the two houses we only use estimates, where house controls are included. The estimate of 0.14 is retrieved from column 1 in table B1 while 0.63 is retrieved from column 2 of table 3b.

	Lower bound	Prefered estimate	Upper bound
P_0	1,680,000	1,680,000	1,680,000
i	0.04	0.04	0.04
T	30	30	30
q_{10}	0.36	0.36	0.36
q_{90}	0.62	0.62	0.62
β	0.14	0.36	0.63
$a_{90} - a_{10}$,DKK	3,494.65	9,039.65	15,701.09
$a_{90} - a_{10}$,USD	513.71	1,328.83	2,308.06

Table 4: Calculation of costs

The table displays assumed values and the calculation of the lower and upper bound of the difference in annuities between two identical houses in SABs with school SES in the 20th and 80th percentile. The exchange rate from DKK to USD is 0.147.

expense of 1,329 USD per year. For perspective, we calculate bounds as shares of income for each percentile in the income distribution of Danish families in 2015. We include families, where at least one adult is between 25 and 35 years old, and where there is at least one child living at home. We use disposable income after taxes and transfers and subtract the annuity value of P_0.[30] The calculated shares are therefore the share of yearly consumption that household would have to forego to buy a house associated with a school with socioeconomically strong students compared to the school with more disadvantaged students. Figure 4 displays the result, where the black lines represent the bounds and the blue solid line represents our preferred estimate.[31] For the households with the highest income the costs are negligible, but for the lower income households, the cost may represent a substantial decrease in disposable income. The household at the 10th percentile will have to give up seven to nine percent of their disposable income. The cost of access to high-SES public schools can, therefore, make up a sizable budget share for low-income families.[32]

Though these calculations are subject to assumptions, they elucidate an important hindrance in ensuring equality of opportunity. Even in a welfare state with extensive

[30]Technically we use the disposable income variable, DISPON, where we add back the rental value of property. We exclude the top 0.1 percent of the household income distribution. We subtract the annuity value of a mortgage of 1,680,000 DKK which equals 97155 DKK. We exclude households with less than 10,000 DKK disposable after subtracting the annuity, which amounts to one percent of households.

[31]The muted lines represent the shares calculated from all estimates of β. Dashed lines represent Difference-in-Difference estimates while the solid lines are from the cross-sectional BDD estimations.

[32]The calculation is of course simplified. We assume that everybody has access to loans at the same interest rate. Furthermore, the fixed costs of obtaining the loan do not enter into the calculation. We also assume that households may get the loan in the first place, which is not necessarily the case. All these assumptions will tend to underestimate the true inequality in access to high-SES primary schools. As we have only exploited variation in prices within municipalities, we have essentially identified after-tax valuations. The back-of-the-envelope calculation therefore implicitly assumes that the two schools are within the same municipality. In 2015, the municipal taxes ranged from 22.5 to 27.8 percent of income, with 50 percent of municipalities within 24.9 and 25.8 percent. The variation in municipal taxes is therefore not great.

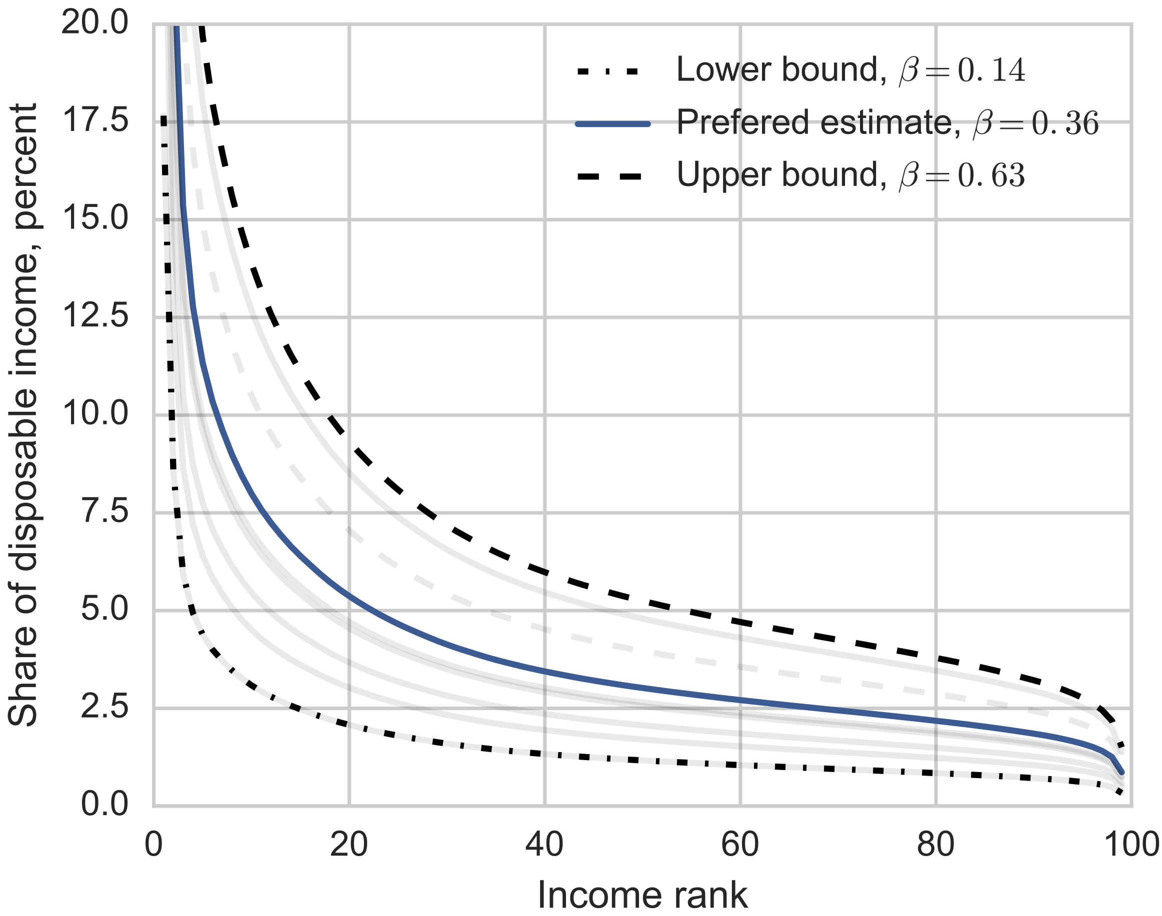

Figure 4: Annuity cost as share of income along the distribution of household income

The calculations are based on the assumptions in Table 4 and the distribution of household disposable income for Danish families in 2015. We restrict the sample to households where at least one of the adult members are between 25 and 35 years old and where there is at least one child. We use the variable DISPON as income concept and subtract the rental value of housing, calculated by Statistics Denmark. We further exclude the top 0.1 percent of the income distribution. We subtract the annuity value of P_0=1,680.000 DKK. We exclude all households with less than 10,000 DKK disposable after the annuity is subtracted, amounting to the lowest one percent. We calculate the curves for all estimates of β where house controls are included. Dashed lines represent estimates are from the Difference-i-Difference approach while solid lines represent estimates from the Boundary-Discontinuity approach.

transfers, poorer households need to give up a substantially larger share of their disposable income to gain access to the same educational services as high-income households. Whether these schools, with stronger peers, would be better for low-income children is a question we do not answer here. In Bjerre-Nielsen and Gandil (2018b) we find supportive evidence that stronger peers are important for low-SES children. This suggests that the cost of housing may be a blockade to achieving the full educational potential of low-SES children.

As long as children are allocated to schools via geographical zoning, we are likely to observe such patterns. A strategy to combat inequality in access is to apply more flexible admission criteria. An example of this is the results found in Machin and Salvanes (2016). The authors show that loosening the geographically defined admission criteria to high schools decreased the capitalization of school quality. If the government wants to increase equality of opportunity, this may be a possible policy. However, increasing the degree of choice may not be sufficient. If the admission systems are complicated, sophisticated households may be able to exploit the system, leaving the less sophisticated households behind. Bjerre-Nielsen and Gandil (2018a) show that this is one of the primary ways high-SES families avoid low-SES peers in primary schools in Denmark. Loosening the admission criteria must, therefore, be done in a transparent way, such that all households may be able to navigate the process. Nevertheless, even if such a policy is implemented, the local nature of educational services will continue to create inequality to some degree if neighborhoods are unequal. Thus, without housing policy, a government is limited in its ability to ensure equal access to education and thereby possibly equality of opportunity.

7 Conclusion

In this paper, we have estimated the sensitivity of house prices to school characteristics. Using both a boundary discontinuity design and changes in school attendance boundaries. we find that prices rise with a socioeconomic index. We find little effect from test scores. We have shown that the implied price differentials between socially strong and weak schools are sizable. Low-income households may, therefore, have to give up a substantial share of consumption to buy their way into socially strong schools. Insofar as strong peers improve child outcomes, the results show that public provision of free education may be unable to ensure equality of opportunity, when children are allocated to schools according to their residential location.

References

Angrist, J. D., Pischke, J.-S. 2008. Mostly Harmless Econometrics: An empiricist's companion. Princeton university press.

Bayer, P., Ferreira, F., McMillan, R. 2007. A Unified Framework for Measuring Preferences for Schools and Neighborhoods. Journal of Political Economy, 115, 588–638.

Benabou, R. 1993. Workings of a city: location, education, and production. The Quarterly Journal of Economics, 108, 619–652.

Benabou, R. 1994. Human capital, inequality, and growth: A local perspective. European Economic Review, 38, 817–826.

Bjerre-Nielsen, A., Gandil, M. H. 2018c. Privacy in spatial data with high resolution and time invariance.

Bjerre-Nielsen, A., Gandil, M. H. 2018a. Defying attendance boundary policies and the limits to combating school segregation.

Bjerre-Nielsen, A., Gandil, M. H. 2018b. Do peers matter? only if you need them (and meet them).

Black, S. E. 1999. Do Better Schools Matter? Parental Valuation of Elementary Education. The Quarterly Journal of Economics, 114, 577–599.

Black, S. E., Machin, S. 2011. Housing Valuations of School Performance. 3, Elsevier B.V. 1st edition, 485–519.

Bogart, W. T., Cromwell, B. A. 2000. How much is a neighborhood school worth? Journal of urban Economics, 47, 280–305.

Durlauf, S. N. 1996. A theory of persistent income inequality. Journal of Economic Growth, 1, 75–93.

Epple, D., Romano, R. E. 1998. Competition between Private and Public Schools, Vouchers, and Peer-Group Effects. The American Economic Review, 88, 33–62.

Fack, G., Grenet, J. 2010. When do better schools raise housing prices? Evidence from Paris public and private schools. Journal of Public Economics, 94, 59–77.

Figlio, D. N., Lucas, M. E. 2004. What's in a grade? school report cards and the housing market. American economic review, 94, 591–604.

Fiva, J. H., Kirkebøen, L. J. 2011. Information shocks and the dynamics of the housing market. The Scandinavian Journal of Economics, 113, 525–552.

Gibbons, S., Machin, S., Silva, O. 2013. Valuing school quality using boundary discontinuities. Journal of Urban Economics, 75, 15–28.

Imbens, G. W., Lemieux, T. 2008. Regression discontinuity designs: A guide to practice. Journal of econometrics, 142, 615–635.

Imberman, S. A., Lovenheim, M. F. 2016. Does the market value value-added? evidence from housing prices after a public release of school and teacher value-added. Journal of Urban Economics, 91, 104–121.

Kain, J. F., Quigley, J. M. 1970. Measuring the value of housing quality. Journal of the American statistical association, 65, 532–548.

Kane, T. J., Riegg, S. K., Staiger, D. O. 2006. School quality, neighborhoods, and housing prices. American Law and Economics Review, 8, 183–212.

Kane, T. J., Staiger, D. O., Samms, G. 2003. School accountability ratings and housing values. Brookings-Wharton papers on urban Affairs, 83–137.

Machin, S., Salvanes, K. G. 2016. Valuing School Quality via a School Choice Reform. Scandinavian Journal of Economics, 118, 3–24.

Ries, J., Somerville, T. 2010. School quality and residential property values: evidence from vancouver rezoning. The Review of Economics and Statistics, 92, 928–944.

Rosen, S. 1974. Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. Journal of Political Economy, 82, 34–55.

Tiebout, C. M. 1956. A Pure Theory of Local Expenditures. Journal of Political Economy, 64, 416–424.

A Additional descriptive statistics

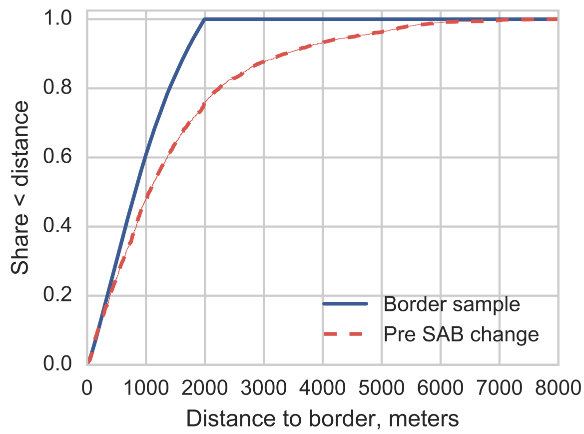

Figure A.1: CDF of distances

The figure displays the cumulative distribution functions for distances to border. For the boundary discontinuity sample distance is censored at 2 kilometres. For those addresses shifted, we only report the distance prior to the change and do not cap the distance. The border used is the border of the district to which the address is eventually transferred.

	(1)		
	School SES	Non_Western share	GPA
School SES	1		
Non-Western share	-0.709***	1	
GPA	0.665***	-0.534***	1

Table A1: Covariance of school characteristics

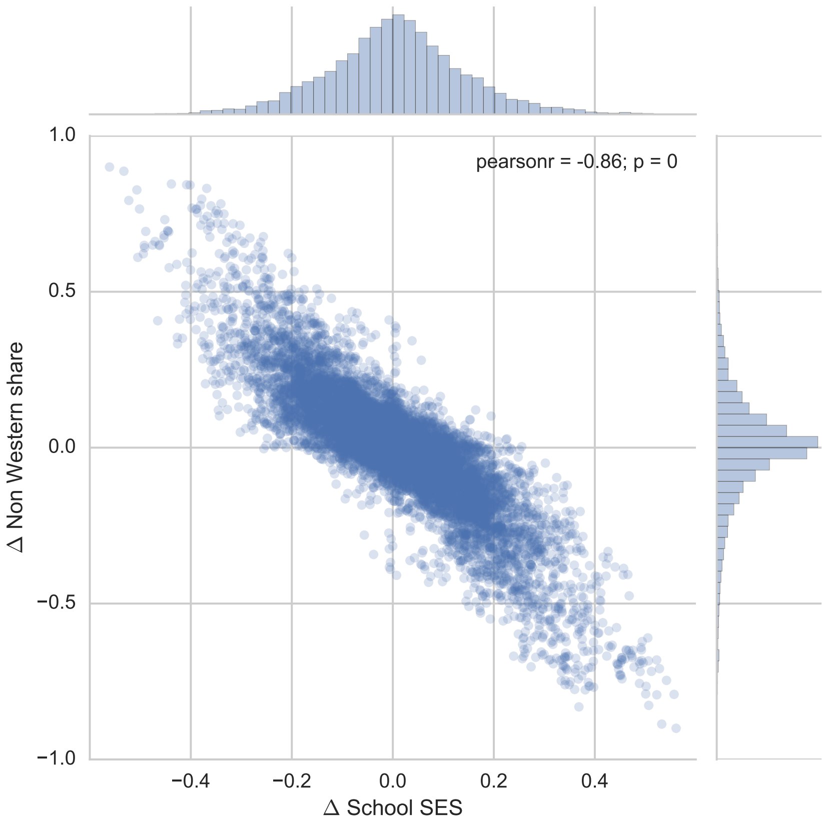

Figure A.2: Joint distribution of border differences in school SES and non-Western share

The figure displays the joint border difference distribution in school SES and non-Western share. Each dot represent a side of the border in a given year. Only border sides with more than ten observations are shown. If all borders were shown, the distribution would be symmetric around zero.

B Static analysis

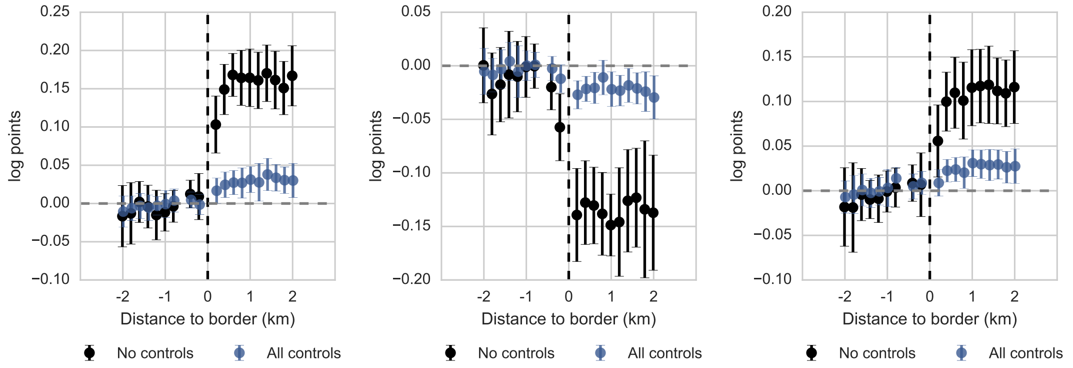

(a) School SES (b) Non-Western share (c) Average GPA

Figure B.1: Simple Boundary Discontinuity Design with confidence intervals

The graphs show the results of a Boundary Discontinuity Design with one school characteristic at a time. We run regressions of discretized distances to the border, where negative distances signify that the address belongs to the side of the border with the lowest value of the measure in question. We include a border-year fixed effect to control for level differences shared by both sides of the border. In black we present the parameters on the binned distance dummies with no additional controls beside the fixed effect. The parameters in blue are from an estimation where we include hyper-local neighborhoods and square meters (including all controls squared). The results are normalized at the 400-meter distance a the left side of the border. Standard error are clustered at municipal level.

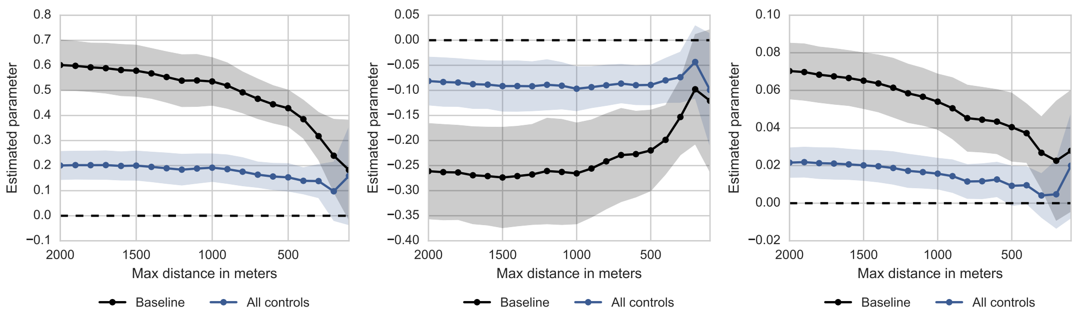

(a) School SES (b) Non-Western share (c) GPA

Figure B.2: Importance of boundaries

The figures present estimations of (9) with different restrictions on distance to the border. Going left to right the restriction becomes more restrictive.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
School SES	0.316***	0.203***	0.139***							0.572***	0.305**	0.246**
	(0.0498)	(0.0480)	(0.0345)							(0.120)	(0.105)	(0.0740)
School NW-share				-0.151***	-0.112**	-0.0733**				0.222*	0.0749	0.0561
				(0.0386)	(0.0342)	(0.0253)				(0.0895)	(0.0670)	(0.0474)
School GPA							0.0265**	0.0129	0.00432	-0.00268	-0.00661	-0.0119
							(0.00983)	(0.00879)	(0.00592)	(0.0105)	(0.0103)	(0.00651)
N: SES		1.560*	0.395		1.583*	0.412		1.535*	0.430		1.480*	0.384
		(0.649)	(0.556)		(0.659)	(0.559)		(0.681)	(0.572)		(0.665)	(0.573)
N: SES squared		-0.273	0.341		-0.281	0.334		-0.271	0.301		-0.231	0.335
		(0.543)	(0.467)		(0.554)	(0.470)		(0.565)	(0.480)		(0.550)	(0.481)
N: NW-share		0.0558	-0.120		0.0563	-0.121		0.0172	-0.157		0.0238	-0.150
		(0.123)	(0.109)		(0.124)	(0.109)		(0.123)	(0.112)		(0.123)	(0.112)
N: NW-share squared		0.149	0.281		0.161	0.290		0.251	0.375		0.232	0.359
		(0.387)	(0.287)		(0.386)	(0.286)		(0.384)	(0.284)		(0.390)	(0.289)
N: Long-cycle Educ (share)		0.516**	0.462**		0.534**	0.475**		0.536*	0.475**		0.522*	0.464**
		(0.184)	(0.142)		(0.185)	(0.144)		(0.207)	(0.168)		(0.204)	(0.164)
N: Long-cycle Educ (share) squared		0.173	-0.475*		0.143	-0.495*		0.119	-0.482		0.131	-0.473
		(0.299)	(0.226)		(0.300)	(0.229)		(0.344)	(0.253)		(0.337)	(0.247)
N: Employment		-1.349*	-0.585		-1.370*	-0.600		-1.176*	-0.477		-1.135	-0.443
		(0.552)	(0.429)		(0.553)	(0.425)		(0.582)	(0.447)		(0.575)	(0.449)
N: Employment squared		0.487	0.183		0.497	0.191		0.381	0.115		0.356	0.0951
		(0.337)	(0.285)		(0.338)	(0.282)		(0.357)	(0.298)		(0.354)	(0.301)
Single Family Home$\times M^2$			0.00302***			0.00302***			0.00299***			0.00299***
			(0.000272)			(0.000271)			(0.000277)			(0.000278)
Terraced house$\times M^2$			0.00509***			0.00509***			0.00508***			0.00508***
			(0.000258)			(0.000259)			(0.000277)			(0.000275)
Apartment$\times M^2$			0.00886***			0.00886***			0.00887***			0.00887***
			(0.000298)			(0.000298)			(0.000309)			(0.000310)
N	56374	56374	56374	56374	56374	56374	52872	52872	52872	52872	52872	52872

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B1: Full regression results

This table present full regression results for estimations of log sales prices on school, house and neighborhood characteristics. The observations are limited to observations within 300m of the boundary. All models are estimated with housetype-border-year fixed effects. Standard errors are clustered at municipal level regardless of year. The prefix "N:" signifies identifies variables calculated as hyperlocal averages.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
GPA	0.0349^{***}	0.0234^{***}	0.0135^{*}	0.0158^{***}	$0.00932^{\dagger}$	0.00432	0.00179	-0.00148	-0.00898	-0.00177	-0.00482	$-0.0119^{\dagger}$
	(0.00555)	(0.00661)	(0.00606)	(0.00418)	(0.00542)	(0.00592)	(0.00603)	(0.00718)	(0.00646)	(0.00509)	(0.00649)	(0.00651)
School SES							0.614^{***}	0.429^{***}	0.405^{***}	0.286^{**}	0.180^{*}	0.246^{**}
							(0.115)	(0.114)	(0.0839)	(0.0928)	(0.0858)	(0.0740)
Non-Western share							0.225^{**}	$0.132^{\dagger}$	0.138^{*}	0.0780	0.0151	0.0561
							(0.0800)	(0.0737)	(0.0549)	(0.0653)	(0.0563)	(0.0474)
Boundary distance	¡1000m	¡500m	¡300m	¡1000m	¡500m	¡300m	¡1000m	¡500m	¡300m	¡1000m	¡500m	¡300m
House controls	X	X	X	X	X	X	X	X	X	X	X	X
Neighborhood controls				X	X	X				X	X	X
N	193537	94495	52872	193537	94495	52872	193537	94495	52872	193537	94495	52872

Standard errors in parentheses
† $p < .1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B2: Marginal GPA effect estimates

The models present estimates from an OLS regression of house prices on school characteristics and a border-type-year fixed effects. Standard errors, clustered at municipal level, are presented below the estimates. We estimate the model under three levels of boundary distances: ¡1000m; ¡500m, and; ¡300m. In some models neighborhood and other school characteristics are included.

C Hedonic regressions and the importance of schools

The formulation OF local neighborhood characteristics as bad controls is the mirror image of the argument presented by Bayer et al. (2007) to estimate preferences for neighbors. In this paper the authors argue that if sorting occurs on aN observable variable then by controlling for this variable one can causally estimate the weight put on other characteristics *over and above* the influence on school characteristics. By conditioning on the boundary fixed effect *and* school quality, we would in line with this argument be able to estimate preferences for local neighborhood characteristics. In this case sorting is seen as a function of school characteristics and not the other way around.

As we initially introduced the hyperlocal neighborhoods merely as controls to ensure exogenous variation in school characteristics we regard the interpretation as susceptible to omitted variable bias. We do not know if the exact characteristics are the characteristics that house buyers care about or merely correlate. We however still find it instructive to inspect how the inclusion of the boundary fixed-effect affects estimates of the neighborhood characteristics. To do this, we estimate hedonic regressions with neighborhood variables and house-level variables, with and without a municipality fixed effect and a boundary-year fixed effect. For the regressions without fixed effects we include dummies for years to remove shared time trends.[33] When we interpret the coefficients on neighborhood characteristics we essentially "flip" our approach compared to the main text; the neighborhood variables are now of interest and the school characteristics serve as controls. We therefore include school characteristcs jointly and abstain from interpreting them.

Column 1 of Table C1 shows the result of a simple hedonic regression. As expected local SES is positively associated with higher prices. Somewhat surprisingly, the non-Western share is significantly positive. However, once we introduce the municipality fixed effects the parameter on the local non-Western share switches signs. The strong reversal most likely reflect urbanization. Non-Western immigrants and descendants tend to cluster in bigger cities where the price level is in general higher. Once we introduce the fixed effects, we are only using within-municipality variation, and the effect from the degree of urbanization therefore disappears. The valuation of neighborhood SES drops by two thirds when we include the municipality fixed effect, implying that different social classes tend to cluster in different municipalities. When we proceed to include border-year fixed effects, the valuations fall further. The parameter on neighborhood-SES is now 0.7, which is less than a fourth of the naive estimate of three in column 1. The parameter on the non-Western share in the neighborhood falls relative to column 2 but maintains the sign. One can interpret this as evidence that households sort on unobservables, even within municipalities. By introducing the fixed effect we control for unobserved amenities

[33]These trends are entirely subsumed into the border-year fixed effects.

and the result is lower values on neighbors than a naive regression would imply, also when municipalities are controlled for.

In column 4 in Table C1 we introduce school level variables and once again omit the fixed effects. Comparing the estimates to column one we see that the neighborhood variables once again move towards zero. The non-Western share is essentially zero. Reintroducing the municipality fixed effect however, the non-Western share rises in magnitude to -0.2 but drops to -0.07 again once we include the border-year fixed effect. The latter estimate is very close to the parameter of column 3 where we included border-year fixed effect but left out school-level characteristics. This conclusion holds for neighborhood-SES as well: Once we control for boundary fixed effects, the parameters on neighborhood characteristics are not sensitive to the inclusion of school characteristics.

The importance of the boundary fixed effects found in Table C1 mirror the findings of Bayer et al. (2007) closely. Taking these results at face value, they imply that traditional estimates of the importance of neighbors are biased due to unobserved amenities. We however once again stress that we regard these results as under-identified and thus are not too comfortable believing in the exact estimates. Nonetheless, these results can serve as validation in the hedonic pricing literature.

	(1)	(2)	(3)	(4)	(5)	(6)
N: SES	3.015***	1.128***	0.738***	1.600***	0.993***	0.722***
	(0.194)	(0.0908)	(0.0487)	(0.103)	(0.0719)	(0.0504)
N: NW-share	0.843***	-0.223*	-0.0634	-0.0416	-0.286**	-0.0626
	(0.176)	(0.0980)	(0.0454)	(0.152)	(0.0860)	(0.0471)
H: Terraced house	-0.333***	-0.344***	-0.329***	-0.348***	-0.346***	-0.329***
	(0.0387)	(0.0485)	(0.0446)	(0.0353)	(0.0444)	(0.0443)
H: Apartment	-0.712***	-0.870***	-0.951***	-0.788***	-0.870***	-0.950***
	(0.0646)	(0.0364)	(0.0398)	(0.0589)	(0.0384)	(0.0398)
H: Single Family Home $\times M^2$	0.00273***	0.00326***	0.00309***	0.00279***	0.00321***	0.00309***
	(0.000230)	(0.000217)	(0.000267)	(0.000221)	(0.000225)	(0.000269)
H: Terraced house $\times M^2$	0.00543***	0.00552***	0.00513***	0.00534***	0.00550***	0.00513***
	(0.000347)	(0.000270)	(0.000303)	(0.000308)	(0.000240)	(0.000299)
H: Apartment $\times M^2$	0.00937***	0.00965***	0.00912***	0.00963***	0.00962***	0.00912***
	(0.000470)	(0.000353)	(0.000360)	(0.000401)	(0.000350)	(0.000360)
S: SES				2.065***	0.438***	0.263**
				(0.116)	(0.121)	(0.0804)
S: NW-share				1.322***	0.293**	0.0709
				(0.0904)	(0.105)	(0.0478)
S: GPA				-0.0448***	0.0100	-0.00963
				(0.0111)	(0.00597)	(0.00646)
Year dummies	X	X		X	X	
Municipality FE		X			X	
Border-year FE			X			X
N	52872	52872	52872	52872	52872	52872

Table C1: Hedonic regressions with and without border-year fixed effects

The models present estimates from an OLS regression of log house prices on neighborhood characteristics ("N"), house characteristics ("H") and school characteristics ("S"). Model 2 and 4 include border-year fixed effects. Standard errors, clustered at municipal level, are presented below the estimates. Only houses within 300 meters of the border enter the regression. Farmhouse is the reference category.

D BDD: effect heterogeneity

One may worry that the effects estimated so far reflect effect heterogeneity. To investigate this we rank the border-year combinations according to the absolute difference between the two neighboring schools. We then run our model with full controls for neighborhood and house characteristics in a running 10 percent sample window. The average difference in log prices as a function of the difference in school-SES is seen in Figure D.1. Figure D.1b we present the estimated differences as a function of the mean difference in school SES on the x-axis. We see that for small differences in school-SES there is no difference in the prices. However, once differences in school-SES exceed 5 points the relationship becomes approximately linear when. This implies that the marginal effect is approximately constant. This can be seen from Figure D.1b where we have calculated the marginal effect. The marginal effect of 0.14)found in column 3 of Table 2) is the average of an approximate (unstable) zero effect for low differences and a stable effect of around 0.2 for SES differences higher than 5 points.[34] Due to the very stable relation for SES-differences for this large region, we conclude that we do not observe a heterogeneity in the price responses according to the level of "treatment", i.e. the difference in school-SES at the border. With constant marginal effects, we can regard house prices as a log linear function of school-SES.

[34]The very unstable relationship in the bottom of the "difference distribution" can be rationalized by once again viewing the model as a Wald-estimator., where a reduced form estimate is rescaled by the first stage to obtain a marginal effect. By construction, the samples in the left most part of the figures have small differences between school SES across the border. This difference constitutes the first stage of an IV. This implies that tiny deviations in the reduced form regression, the nominator of the Wald estimator, will cause the the ratio between the reduced form and the first stage to explode in magnitude.

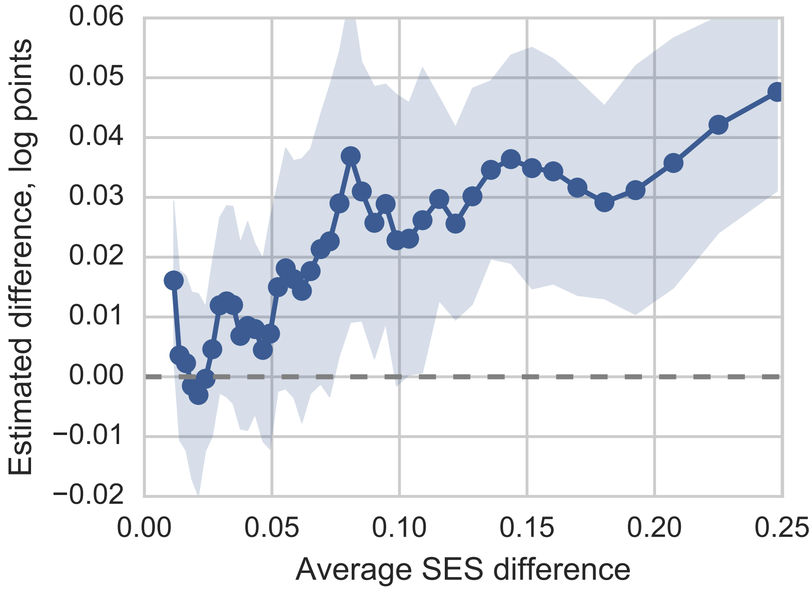

(a) Difference at border

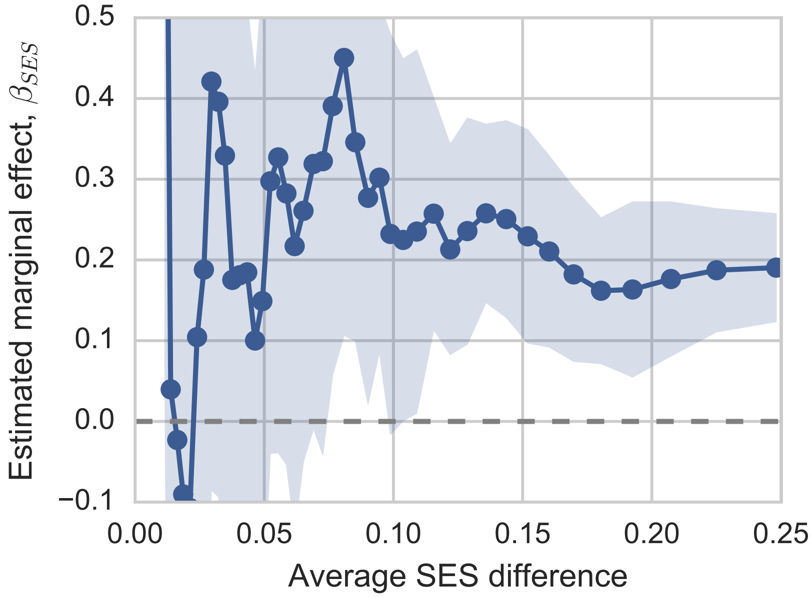

(b) Marginal effect

Figure D.1: Simple Boundary Discontinuity Design

The graphs show the results of a Boundary Discontinuity Design with SES as a measure of schools. In Figure D.1a we regress log prices on a dummy for being on the right side of a district border as well as controls for house type interacted with square meters as well as hyperlocal neighborhood measures (and squares). We include a border-type-year fixed effect to control for level differences shared by the same types of houses on both sides of the border in a given year. We compute a parameter for a sliding 10 percent window according to the ranked border regions. In Figure D.1b we compute the marginal effect by exchanging the right side dummy for the level of school SES. 95-percent confidence bands are displayed as the shaded area. Standard errors are clustered at municipal level.