

FAITH PRIMARY SCHOOLS: BETTER SCHOOLS OR BETTER PUPILS?

Stephen Gibbons*, Olmo Silva**

December 2007

Author for correspondence:

Olmo Silva,
Department of Geography and Environment,
London School of Economics,
Houghton Street, WC2A 2AE, London, UK.
Tel.: 0044 20 7852 3516
Fax: 0044 20 7955 7595
Email: o.silva@lse.ac.uk

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

**Department of Geography and Environment and Centre for Economic Performance, London School of Economics; and IZA, Bonn.

Acknowledgements: We would like to thank Joseph Altonji, Steve Machin, Sandra McNally, Steve Pischke, Anne West and participants at the informal CEE Meeting, IZA/SOLE Annual Meeting 2006, CEP-LSE Labour Market Workshop, EEEPE Summer School 2006 and DfES internal seminars, for their helpful comments. We also would like to thank Anushri Bansal for excellent research assistance. We are responsible for any errors or omissions.

FAITH PRIMARY SCHOOLS: BETTER SCHOOLS OR BETTER PUPILS?

Abstract

We provide estimates for the effect of attending a Faith school on educational achievement using a census of primary school pupils in England. We argue that there are no credible instruments for Faith school attendance in this context. Instead, we partially control for selection into religious schooling by tracking pupils over time and comparing attainments of students who exhibit different levels of commitment to religious education through their choice of secondary school and residence. Using this approach, we find only a small advantage from Faith primary schooling, worth about 1 percentile on age-11 test scores. Moreover, this is linked to autonomous admissions and governance arrangements, and not to religious character of the schools. We then go on to show that our estimates vary substantially across pupil sub-groups that exhibit different levels of sorting on observable characteristics into Faith schooling, and provide bounds on what the ‘Faith school effect’ would be in the absence of sorting and selection. Pupils with a high degree of observable-sorting into Faith schools have an age-11 test score advantage of up to 2.7 percentiles. On the other hand, pupils showing a very low degree of sorting on observables have zero or negative gains. It appears that most of the apparent advantage of Faith school education in England can be explained by differences between the pupils who attend these schools and those who do not.

Key words: Faith school; Primary schools; Pupil achievement.

JEL Classification: I20; J24; Z12.

1. Introduction

Quality of childhood schooling is increasingly seen as important for life chances, because adults' success in the labour market is closely linked to early educational attainments (Heckman, 2000). However, economic and educational research has had little success in identifying resource-based interventions that are effective in raising school standards (Hanushek, 2003). Government policy in many countries now favours reforms based on incentives, governance, increased choice and competition¹. In England, this idea has become linked with the expansion of the state-funded Faith-school sector because it symbolises choice and diversity in the education system, it embodies the kind of practice in governance that policy makers wish to promote, and – crucially – because it is claimed it offers higher educational standards (DfES White Paper, 2005; DfES Education Bill, 2006). This case is, however, a difficult one to assess, because pupils that choose and get chosen by Faith schools differ from the population of pupils in ways that are correlated with their educational achievement.

Surprisingly, there is almost no evidence on the issue for England that makes any serious attempts to separate out the 'causal' effect of attending a Faith school on educational achievement from pure selection. Much more work has been done on this in the US – in relation to Catholic schools in particular. However, one important difference between the US and English settings is that, while US Faith schools are private institutions, in England they are state funded and part of the mainstream state system. This makes England a more attractive setting for studying the specific effects of religious affiliation and ethos on academic achievement, as distinct from the effects of a private education. In addition, it makes this country an interesting framework for assessing the effectiveness of policies aimed at improving state-funded education through the expansion of choice

¹ See LeGrand (1991) and (1993), and Machin and Vignoles (2005) for a review of the English experience, and Hoxby (2004) for an analysis of US based evidence.

and alternative institutional arrangements. Nevertheless, the problem of non-random sorting of pupils into Faith schools remains.

In this paper we present new evidence on this issue. We provide insights into the extent to which the gap between Faith and Secular schools in terms of student outputs is driven by better schooling in the Faith sector, or by the fact that Faith schools admit ‘better’ pupils to start with. Previous (mainly US) research that has tried to tackle this topic has typically made use of instruments for Faith school attendance such as family religion, neighbours’ religion and other characteristics of place of residence. However, we agree with claims that these instruments are not credible when the point of the exercise is to purge estimates of family background and ability-related effects (see Altonji et al. 2002, 2005). We argue that these instruments are generally inappropriate, because family religion is correlated with other background characteristics (explicitly so if Faith-school attendance affects pupil outcomes), and because families choose where to live for reasons often related to the school they wish their children to attend².

Instead, we draw on a national census of pupils which allows us to compare standard national test scores for students who attended Faith schools in the primary phase with similar students in the Secular (non-Faith) primary sector. Using these data we can control carefully for prior attainments, family characteristics, and – importantly – place of residence. Indeed, in contrast to previous literature, we argue that consideration of the process of residential choice means that it is better to *control* for precise residential location, rather than use it to derive instruments. To this end, we exploit the geographical detail in our data set to compare outcomes for primary school pupils who live in the same postcode (10 or so contiguous housing units), but attend different schools. In addition, our methods take advantage of the fact that we can observe pupils at two phases of their education (primary and secondary). This feature of the data helps us to consider the influence of

² Indeed, it is likely to be impossible to find instruments that induce random assignment to Faith schools, because school and residential choice is always subject to personal preference and any random assignment without compulsion could be undone by individual action.

unobservable characteristics – especially parental commitment to religious education and school-side selection of ‘suitable’ pupils – that are revealed in subsequent choice of secondary school.

Even then, without some form of random assignment, we do not believe that we can draw firm conclusions about a unique parameter that characterises the average causal influence of Faith schools. Instead, we present estimates under a range of different specifications in which we compare sub-groups of pupils whose families make either very similar or very dissimilar educational choices: some never attend a Faith school, some attend a Faith primary school but not a Faith secondary school or vice-versa, and some attend Faith schools in both phases. We argue that these patterns of choice are revealing about unobserved family preferences and characteristics, as well as school-side admission policy. Comparison of the Faith school test score gap for these sub-groups, relative to their degree of sorting on observable characteristics into religious education, allows us to quantify how much of the Faith school advantage can be explained by individual characteristics and bound the likely magnitude of what the causal impact of Faith schooling would be in the absence of selection and sorting.

To preview our results, we find that most, if not all, of the observed educational advantage of Faith schools is due to non-random selection of higher ability and better-background children into schools that have autonomous admissions and governance arrangements – including those in the Faith sector. Attending a school of this type between ages 7 and 11 raises a child’s attainment by between zero and 2.7 percentiles in the distribution of pupil test scores at age 11, conditional on age 7 attainment, with our central estimate being below one percentile. Even then, this small advantage is not specifically linked to religious affiliation: Secular schools, with similar autonomous admissions and governance arrangements, show a comparable performance advantage, whilst Faith schools under closer control by the local education authorities show none.

The paper has the following structure. The next section critically reviews some of the literature on the topic and its methods. In Section 3 we explain the different types of schools that exist in the

English school system, and the data that we will be using. Section 4 sets out our empirical methods in more detail, and Section 5 presents and discusses the results that arise from these approaches. Section 6 concludes.

2. A review of the methods used in previous studies

By far the bulk of existing academic work on the impact of Faith schools on education originates in the US, and has focussed on Catholic schools – largely springing from the influential work of Coleman (1982)³. Most of the subsequent research there finds that attendance at a Catholic school raises students' graduation rates and test scores, though there is variation across different demographic and geographical groups, and across subject areas.

For Britain, evidence on the performance benefits of Faith schooling is fairly limited. Schagen et al. (2002) show that pupils in Faith secondary schools progress faster in English (but not in Maths and Science), and also seem to pass more subjects overall in their age-16 exams. Benton et al. (2003) report that Faith secondary schools are associated with faster grade progression between age 11 and 13, and age 13 and 16, but this is confined to schools affiliated with non-mainstream Christian (i.e. not Catholic or Church of England) and Jewish denominations. Finally, in a study limited to two London boroughs with only 7 religious schools, Prais (2005) finds quite strong Faith school advantages in Maths, particularly amongst weakest pupils. However, neither of these studies takes any steps to control for pupil background or otherwise deal with selection on unobservable characteristics that influence educational progress.

Our discussion in this section will mainly focus on whether there is anything we can learn from the methodological approaches used in the US and international literature, rather than the results per-se. As noted above, an important difference between the US and England is that Faith schools in the US – mainly Catholic – are private-sector schools whereas Faith schools in England –

³ See Dronkers (2004) for a review of the European evidence on educational and behavioural effects of attending a Faith school, and a detailed analysis of the historical and institutional determinants of religious education in Europe.

mainly Catholic and Church of England – are part of the state-school system. Nevertheless, the key issue that has taxed US researchers investigating the effects of Catholic schooling on achievement is the same one that we will need to confront here: there is clearly non-random sorting of pupils into Faith schools, such that religious school attendance is correlated with unobserved pupil and family characteristics that are educationally advantageous. Almost all approaches try to find an explicit source of random variation in the probability of Faith school attendance that is otherwise uncorrelated with educational attainment and can be used as an instrument. Disappointingly, many of the instrument choices do not seem credible on deeper reflection, and the evidence in Altonji et al. (2002) and (2005) is not supportive of any of those that are commonly used.

The first typical instrument is family religion, on the basis that being Catholic is a strong determinant of attendance at a Catholic school. This approach is used, for example, by Noell (1982), Evans and Schwab (1995) and in part of Neal (1997). However, opinion seems divided on whether family religion is related to educational outcomes other than through Catholic school attendance. On a priori grounds, it seems most likely that religious beliefs influence all sorts of family attitudes and economic outcomes⁴. In fact, most recent studies show that the range of family outcomes that are correlated with religiosity may be much wider: Gruber (2005) finds evidence that religious participation is correlated with income, lower rates of disability, and with more marriage and less divorce. Another strand of literature links religious beliefs to economic growth, ranging from sociological work such as Weber (1905) to macroeconomic studies such as Barro and McCleary (2003). In a similar vein, Guiso et al. (2003) use international micro data to show a link between religion and social attitudes that are conducive to positive economic outcomes. On balance, we are not convinced that family religion is a useful instrument for Faith school attendance.

Another approach has been to use instruments that try measure the local ‘supply’ of Faith schooling. Neal (1997) uses both the geographic density of Catholic schools and the number of

⁴ This is especially true if attendance at a Catholic school does raise attainments, the probability of graduation, future earnings and hence family resources in future generations – which is what most of these schooling studies imply.

Catholics as proportion of the local population as an instrument, with the justification that Catholic schools charge lower fees in predominantly Catholic areas. Similar ideas are applied in Grogger and Neal (2000) and in Figlio and Stone (1999), who go further and include various geographical and area-demographic variables in their instrument set. We find these ideas unconvincing theoretically: observational evidence on the proximity of place of residence to Faith schools is most likely related to family's preferences over schooling because the decision to live near a religious school is based on the intention to attend that school. The role of schools in housing choices is evident in the voluminous empirical literature on the influence of schools on housing demand (Black, 1999; Gibbons and Machin, 2006; Kain, Staiger and Reigg, 2005). Similarly, local demographic characteristics, such as the local proportion of Catholics, must be related to pupil's family background because his or her family has made choices to live in this type of community. In particular, one reason why a locality may have high concentrations of, say, Catholic families is precisely because these families want their children to attend a local Catholic school. In this case it is the spatial correlation in family preferences over schooling that generates an association between Catholic school attendance and local Catholic density. Other approaches have tried to use interactions of these instruments, whilst controlling for their levels (e.g. Sander, 1996), though the theoretical basis for this is uncertain, and the testing in Altonji et al. (2002) suggests it is not satisfactory.

Given the weaknesses in the IV approach, some have tried other methods. Jepsen (2003) uses value-added models to control for pupil background characteristics and finds no impact of Catholic schools on test scores⁵. In a different approach, Altonji et al. (2005) infer the degree of selection bias in the Catholic school effect from the extent of selection on observable pupil characteristics, and conclude that whilst there is an impact on high school and college graduation rates, there is no influence on test scores. The approach we will follow here is closer to these more recent US studies,

⁵ For the sample of fourth graders, the author also uses school-median test scores of a younger cohort (first-graders) to control for school-level selection, and comes to similar conclusions.

but we have the advantage of a dataset on the population of students in England, containing information on pupil's prior attainments, demographics and details on precise geographical location. Before discussing our methods, we outline the institutional context for Faith schools in England, and the details of the dataset we will use.

3. Institutional context and data

3.1. School types and governance

Primary schools in the state-sector in England fall into a number of different categories, and differ in terms of the way they are governed, the ownership of the school buildings, and who controls pupil admissions. The key differences between these school types – Community, Foundation, Voluntary Aided and Voluntary Controlled – are set out in Table 1. In addition there is a small private, fee-paying sector, which we do not consider here⁶. All state schools are funded largely by central government, through Local Authorities (formerly Local Education Authorities or LEAs) that are responsible for schools in their geographical domain. Schools, other than Community schools, are usually linked to a Faith or other charitable organisation ('foundations'), or have 'partnerships' with one or more local organisations.

All schools are run by a Governing Body composed of members elected from amongst parents and staff (Parent Governors and Staff Governors), appointed by the LEA (LEA Governors), appointed by the Faith or charitable organization that owns the school premises (where relevant – Foundation/Partnership Governors), and appointed from the community (e.g. local businesses) by the Governing Body. The Governing Body sets the strategic direction of the school, draws up school policies, sets targets and monitors performance, although day-to-day running is down to the head-teacher and his or her leadership team. The constitution of the Governing Body is important because it determines how much influence various 'stakeholders' have in the way the school is run

⁶ Private schools educate around 6-7% of pupils in England as a whole.

– in particular, the balance between influence by the LEA and influence by the Faith or charitable foundation/partnership. Moreover, in Voluntary Aided and Foundation schools the Governing body is responsible for admissions and has some flexibility in deciding which pupils will be enrolled when the school is oversubscribed⁷. In all other cases, the LEA is the admissions authority. Over the period relevant to our empirical work, applications for admission to Voluntary Aided and Foundation schools were made directly to the school, though this has now changed and the Local Authority coordinates all admissions.

In principle, admissions to all types of school is on the basis of parental preference, but in practice demand often outstrips the supply of places in popular schools, and these are rationed on the basis of various school-specific oversubscription criteria (which must conform to the current Department of Education and Skills Code of Practice on Admissions). For Faith-based Voluntary Aided schools one of the overriding considerations is attendance at a local church, or a reference from a local minister. For Secular schools, instead, the distance between a child's home and school is one of the key considerations. Because of these criteria – alongside the constraints of travel costs – residential choice and school choice decisions are very closely linked. Even so, most households will have a choice of more than one school, and a choice between Secular and Faith schooling, available from where they live.

Schools are further differentiated according to who is classed as the employer – either the LEA or the Governing Body. Although it is the Governing body which is responsible for making staff appointments at all schools – including the appointment of the head-teacher – in Voluntary Aided schools (and to a lesser extent in Foundation schools), the charitable organizations/partnerships have a strong say *via* their representatives in the Governing body. These distinctions are important when we consider the role of Faith schools, because Faith schools are often different in other ways than just religious affiliation (summarized in Table 1).

⁷ Although, according to the national Code of Practice on Admissions, primary schools should not use explicit selection by aptitude.

Given our emphasis on religious education, we re-arrange schools into four types that we feel best characterise their religious affiliation and governance/admissions arrangements. The breakdown is as follows:

- *Secular, non-autonomous*: includes schools that have no religious affiliation and are Community or Voluntary Controlled.
- *Secular, autonomous*: includes schools that have no religious affiliation but are Foundation or Voluntary Aided.
- *Faith, non-autonomous*: includes schools that have a religious affiliation and are Voluntary Controlled.
- *Faith, autonomous*: includes schools that have a religious affiliation and which are Foundation or Voluntary Aided.

The goal of our empirical work will be to explore differences in attainment of pupils in these schools in the primary phase, and to consider to what extent any differences can be attributed to their religious character⁸. First, however, we need to briefly explain the way attainment is assessed in English primary schools, and describe the data we will bring to bear on this question.

3.2. *National curriculum and assessment*

Compulsory education in England is organised into five stages referred to as Key Stages. In the primary phase, pupils enter school at age 4-5 in the Foundation Stage (not to be confused with Foundation schools) and then move on to Key Stage 1, spanning ages 5-6 and 6-7. At age 7-8 pupils move to Key Stage 2, sometimes – but not usually – with a change of school⁹. At the end of Key

⁸ Almost all Faith schools in England are either Church of England or Catholic school. Only a minority of schools (enrolling less than 1% of pupils) is associated with other Faiths (e.g. Jewish or Muslim). Further, note that Catholic schools account for 53% of the Faith, autonomous sector, while nearly all schools in the Faith, non-autonomous sector are affiliated to the Church of England. We will briefly discuss the link between pupil attainments and denominational differences later in the paper.

⁹ In few cases there are separate Infants and Junior schools (covering Key Stage 1 and 2 respectively) and a few LEAs still operate a Middle School system (bridging the primary and secondary phases); we do not consider these schools here.

Stage 2, when they are 10-11, children leave the primary phase and go on to secondary school where they progress through Key Stage 3 and 4. At the end of each Key Stage, pupils are assessed on the basis of standard national tests, and progress through the phases is measured in terms of Key Stage Levels, ranging between W (working towards Level 1) and Level 5+ in the primary phase. A point system can also be applied to convert these levels into scores that are intended to represent about one term's (10-12 weeks) progress.

3.3. The data

The UK's Department for Education and Skills (DfES) collects various data on school and pupils centrally, because the pupil assessment system is used to publish school performance tables and because information on pupil numbers and characteristics are necessary for administrative purposes – in particular to determine funding. A National Pupil Database (NPD) holds information on each pupil's assessment record in the Key Stage tests throughout their school career.

Since 2002, the DfES has also collected information on pupil's school, gender, age, ethnicity, language skills, any special educational needs or disabilities, entitlement to free school meals and various other pieces of information via the Pupil Level Annual Schools Census (PLASC), which is incorporated into the test-score information in the NPD. Importantly, PLASC has information on postcode of residence: a postcode is typically 10 contiguous housing units, which allows us to control very carefully for residential location.

The NPD thus provides a large and detailed dataset on pupils' characteristics and their test histories, with details on the Levels reached in the core subject areas – Maths, English, Science (Science is only tested after Key Stage 1) – and, for Key Stage 2 and beyond, the raw scores in the component tests. In our analysis, we use information on two cohorts: those aged 10-11 and sitting their Key Stage 2 tests in 2002 and 2003, who took their Key Stage 1 tests in 1998 and 1999 respectively. We can also infer to which school these pupils are assigned when they move on to

secondary school in 2003 and 2004. Various other data sources can be merged in at school level – in particular each school’s religious affiliation and the institutional types described above in Section 3.1 – which are available from the DfES ‘Edubase’ system.

We will use this large and complex combined data set – which gives us information on around 1 million pupils in over 14,000 primary schools in England – to estimate the influence of Faith schools on pupil progress through Key Stage 2 (between ages 7 and 11). In the next section we set out the empirical model more precisely.

4. Empirical Model

As discussed above (see Section 2), measurement of the effectiveness of Faith schools presents a challenge because families with a preference for schools with a religious tradition and ethos may, on average, have characteristics that influence academic progress in their children. In addition, in the English context, many Faith schools at both primary and secondary level had until recently (and for the period under analysis) much greater control over their own pupil admissions than did most Secular schools. In fact, Faith schools (and other schools classed as Voluntary Aided or Foundation; see Section 3.1) were allowed to interview or otherwise screen families – ostensibly to determine their religious or other ethical convictions. However, it has long been suspected that this lead to some form of covert selection based on parental and pupil characteristics that are correlated with pupil ability¹⁰.

These forms of school-side and family-side selection together mean that pupils are sorted into schools along lines of ability and family background, with higher strata potentially over-represented in Faith schools. As detailed above (Section 2), we do not believe there are any credible instruments for Faith school attendance – at least not in the current English education setting. But we do have a wealth of information on the residential location and school attendance history of pupils in our

¹⁰ West (2005) and West and Hind (2003) present detailed qualitative evidence on this issue, while Allen (2006) provides some statistical analysis.

sample which, we argue, we can turn to our advantage. We use this information in four ways to improve and bound our estimates. Firstly, we carefully control for pupil prior attainment and estimate flexible models of pupil educational value-added. Secondly, we compare pupils living in close proximity, in order to condition out income, family background and preference-related factors that are linked to residential choice. Thirdly, we condition on future school choices which, we argue, allows us to compare pupils (and their families) who exhibit similar propensities for a Faith-based education. Lastly, we explore the sensitivity of our findings to unobservable factors by comparing our estimates on various sub-samples of pupils who exhibit different school choice patterns and different degrees of sorting along observable dimensions into religious schooling. In the remainder of this section, we explain our reasoning in more detail.

The basic model we will estimate is a standard pupil-level value-added model of educational attainment, which measures the association of school attendance and other characteristics with progress at school between the ages of 7 and 11. In our two-period empirical setup, attainment of pupil i in school j at stage two (h_{ij2}) builds on prior attainment at stage one (h_{ij1}), and is modified by school-type factors (β_j , a school effect that is identical for different pupils in school type j), observable personal/family characteristics, x'_{ij} , and unobserved pupil/family characteristics, η_i :

$$h_{ij2} = \beta_j + x'_{ij}\gamma + g(h_{ij1}) + \eta_i + \varepsilon_{ij} \quad (1)$$

The key empirical problem is that family-side selection of schools, and school-side selection of pupils before stage two, imply that unobserved pupil/family characteristics, η_i , that influence the rate of progress between stage one and stage two, also influence school choice probabilities, so that $\Pr(j = k) = F(x_{ij}, h_{ij1}, \eta_i)$ and $E[\eta_i | h_{ij1}, x_{ij}, j = k] \neq E[\eta_i | h_{ij1}, x_{ij}, j \neq k]$. Estimates of β_j that do not control for $E[\eta_i | j]$ are biased estimates of the expected impact of Faith school attendance. Most of the research on Faith school effects has had to make-do without information on prior attainment. Since we regard selection on ability or achievement as particularly important, we instead allow

these to enter our model in a very general way as indicated by $g(h_{ij1})$. However, as shown by Manning and Pischke (2006), simply conditioning on prior attainment and observable pupil characteristics is an inadequate way to control for pupil selection into different school types. Therefore we need to take further steps to consider the likely impact that η_i has on our regression estimates.

Previous literature on Faith schooling has used instruments – typically geographical – to deal with the problem of selection on unobservables. However, basic theories of urban economics suggest that choice of place of residence is revealing of the benefits of different locations to different types of people: households sort into communities according to the benefit they can derive from local amenities and according to the income they have available to pay for housing or taxes (see Tiebout, 1956). As discussed in Section 3.1, residential choice and school choice decisions are closely linked, and home location will determine the range of feasible schooling options. Given this, it seems to us that one of the most fundamental things to do when looking for evidence of performance advantages in Faith schools is to control for place of residence, rather than use it to predict Faith school attendance as common in instrumental variable approaches (see the discussion in Section 2). The geographical detail and density of pupils in our data means we are able to do this effectively, because we can identify groups of pupils who live in the same postcode (10-12 contiguous housing units), but attend different schools. We therefore allow for postcode fixed effects in our models and compare pupils who are close residential neighbours.

Clearly, this is not enough to eliminate selection into Faith schooling on the basis of unobservable factors in (1), because school choice is linked to preferences even amongst close neighbours with similar resources and residential preferences. Moreover, if schools have some scope to screen pupils, the probability of admission to a Faith school may be determined by

educationally relevant attributes that these schools can observe, while they are unobservable to us¹¹. However, the range of school types in primary schooling is replicated at the secondary phase and we argue that choice amongst these school types at age-12 (which we can observe in our data) is revealing both about preferences of families regarding Faith schooling and their ‘suitability’ in the eyes of Faith school admissions authorities. Unobservable factors that are correlated with choice of primary school are likely to be closely correlated with choice of secondary school, under the assumption that ‘selection’ into Faith schools occurs along similar lines in the primary and secondary phases¹². On the parental/pupil side, this requires that higher ability or better background pupils are not more or less likely to apply to Faith primary schools than Faith secondary schools, and that the effort required to access Faith education is similar at both phases. This last assumption is reasonable, since the criterion for admission usually boils down to attending a local church for a few months. On the school side, we need screening to operate on the basis of pupil and family characteristics, and not explicitly on academic *progress*. This assumption is also quite plausible, because neither primary nor secondary schools have any measure of a pupil’s prior academic progress at the time they admit them¹³. Moreover, it is in line with the evidence in West (2005) and West and Hind (2003). With these considerations in mind, we argue that pupils who live in the same place and attend the same type of secondary school, or attend the very same secondary school, are likely to be much better matched in terms of unobservable characteristics and educational preferences than are pupils who live in different places and attend different secondary schools, or different secondary school types. So, by including postcode-of residence fixed effects and

¹¹ In the US private Catholic school setting, these factors are theoretically related to the benefits of choosing a Faith school, since attendance at a private school rather than a public school imposes financial costs. In England, conditional on place of residence, admission to a state Faith school does not incur high additional costs relative to a non-Faith school. The only likely cost is the effort of demonstrating some religious commitment through church attendance.

¹² Notice that the use future information to control for unobservables is similar to Grogger (1995) on the effect of arrests on labour market outcomes. The author compares the earning of a sample of individual arrested in 1984 or earlier – ‘treated’ group – to those of a sample of individuals whose first arrest occurred after 1984 – ‘control’ group.

¹³ Primary schools admit pupils before any testing has taken place, and pupils apply and receive admission offers to secondary schools before they have taken their Key Stage 2 tests.

secondary school fixed effects, or secondary school type fixed effects, in the primary phase value-added model of equation (1), we can in part condition out $E[\eta_i | j]$ in a non-parametric way.

Even then, we believe that this strategy alone is not enough to allow us to draw firm conclusions about a unique parameter identifying the causal influence of Faith schooling. Close neighbours attending the very same secondary school, but originating from different primary schools, might still differ along unobservable lines. We therefore go one step further by comparing various pupil sub-groups in ways that allow us to gauge whether any link between test scores and Faith school attendance can be explained by selection on unobservable pupil characteristics.

Drawing on Altonji et al. (2005), we argue that the degree of selection into Faith schools along observable dimensions h_{ij1} and x'_{ij} in (1) provides some guidance regarding the amount of selection in terms of unobservable factors η_i . Hence, our approach is based on contrasting groups who exhibit different transition patterns across school types when moving from primary to secondary education, and who exhibit distinctly different patterns in terms of sorting on observable characteristics into Faith schooling. Firstly, we compare students from families that commit to religious schooling over both educational phases with pupils who attend non-Faith schools in both phases. Whatever unobservable and observable characteristics lead different types of individuals to choose or be chosen by Faith schools, it is likely - and we will show this is the case - that there are much bigger differences between pupils with 'Faith-Faith' and 'Secular-Secular' school choice patterns (i.e. the 'stayers'), than between pupils in Faith primary schools (irrespective of their secondary school choice) and the rest of the population. Secondly, we compare pupils who only attend Faith schools in the secondary phase with pupils who attend Faith schools only in the primary phase. There is less reason to expect these two 'switcher' groups to differ widely in terms of education-related characteristics, and indeed we show that there is far less sorting on observable lines across Faith-Secular and Secular-Faith 'switcher' groups. Exploiting the intuitions in Altonji

et al. (2005), we claim that by comparing the changes in our estimates $\hat{\beta}_j$, with the changes in sorting on observable characteristics $E[h_{ij1} | j]$ and $E[x_{ij} | j]$ as we move from the whole population of pupils to the subgroups of ‘stayers’ and ‘switchers’, provides a strong indication of the extent to which differences in $\hat{\beta}_j$ arise through differences in selection on unobservable factors $E[\eta_i | j]$. In turn, this allows us to provide bounds on what the ‘Faith school effect’ β_j would be in the absence of sorting and selection on unobservables.

To assess this strategy more formally, we use a framework similar to Altonji et al. (2005), and compare selection on observables and unobservables into Faith primary schooling in our sub-groups of ‘stayers’ and ‘switchers’. The foundation of this analysis is to estimate a Heckman selection model of the type:

$$\begin{aligned} h_{ij2} &= \beta_j d_{ij} + x'_{ij} \gamma + g_1(h_{ij1}) + \eta_i + \varepsilon_{ij2} \\ d_{ij} &= x'_{ij} \lambda + g_2(h_{ij1}) + \rho \eta_i + v_{ij1} \end{aligned} \quad (2)$$

In which d_{ij} indicates Faith primary school attendance¹⁴. The parameter ρ measures the correlation between unobservables in the Faith school selection equation and in the value-added equation. This parameter captures the degree of selection on unobservables into Faith schooling.

Although the parameter ρ is not properly identified without exclusion restrictions (i.e. a valid instrument), it can be constrained to predefined values in system (2). By changing values of ρ , we can explore the sensitivity of $\hat{\beta}_j$ to different assumptions about the degree of selection on unobservables into Faith schooling and find the value of ρ that is necessary to drive estimates of $\hat{\beta}_j$ to zero. Additionally, we can estimate the degree of selection on observables within this framework for the sub-samples of ‘stayers’ and ‘switchers’ using the correlation between the

¹⁴ In the empirical work we restrict attention to attendance at Faith, autonomous primary schools.

predictions $x'_{ij}\hat{\gamma} + \hat{g}_1(h_{ij1})$ and $x'_{ij}\hat{\lambda} + \hat{g}_2(h_{ij1})$ ¹⁵. These correlations provide a guide to the validity of our claims about the differences in selection on observables into Faith primary schooling between our ‘switchers’ and ‘stayers’ pupil groups. Moreover, by comparison with values of ρ that drive the estimates of Faith school effect to zero, they provide some indication of the likelihood that positive estimates of $\hat{\beta}_j$ occur because of selection on unobservables.

In conclusion, we argue that the patterns of school choice by ‘stayers’ and ‘switchers’ are revealing about unobserved family preferences and characteristics, as well as school-side selection of ‘suitable’ students, and that comparison of the association between Faith school attendance and attainment amongst these different groups allows us to gauge to the magnitude of any causal impacts from the Faith aspects of schooling. Despite being unable to pin down a unique parameter for the average causal effect of religious education, we believe that much can be learnt about the relative role of selection and institutional differences by comparison of the relationship between Faith-primary attendance and attainment across different groups.

5. Results and discussion

5.1. Descriptive statistics

The basic facts about the association of pupil age-11 attainments and the type of primary school attended are summarised in Table 2. The school categories were explained in Section 3 above. The table shows the means and standard deviations of pupil test scores in standard age-11 tests, where the raw test scores are converted into percentiles. Notice that in all the empirical analysis that follows we will work with an average of the pupil’s percentile in the Maths and English distribution because we found no interesting differences between these two subjects¹⁶. Summary statistics are

¹⁵ These are obtained from unconstrained versions of system (2), where identification is (partially) achieved using the non-linearities that characterize the Heckman-two step models. Specifically, we use a probit specification to model selection into Faith, autonomous primary schooling.

¹⁶ Results for Maths and English separately are not reported for space reasons; they are available from the authors.

shown in Row 1 for the whole sample, and then split by broad school type. The figures show the key feature that we wish to analyse: pupils emerging from primary schools that are classified as Faith schools under our definitions (see Section 3) have higher levels of attainment than those emerging from Secular schools. The difference is about 4.75 percentiles in the pupil test score distribution.

Splitting this gap into the finer school classifications defined above, we see that the apparent ‘Faith school’ effect in Row 2 is more specifically associated with Faith schools that we class as autonomous – which means, amongst other things, that they operate admissions policies that are potentially covertly selective. Secular schools with comparable institutional arrangements similarly show better average performance than other Secular, non-autonomous schools. At the time covered by this research, these schools required parents to apply directly to the institution, which then reviewed the applications and was allowed to interview families prior to admission¹⁷. The question we want to address here is to what extent any Faith or autonomous school advantage is a by-product of differences in background characteristics between pupils who enter these schools and those who do not.

5.2. Differences in pupil background and initial attainment

Firstly, we show that there are indeed important and significant differences between school types in terms of the observable characteristics of pupils at the beginning of the age 7-11 phase. Table 2 reports overall means and standard deviations of age-7 attainment and background characteristics in Row 1, and results from regressions of these characteristics on school-type dummies (with Secular, non-autonomous schools as the baseline) in Rows 2-5.

It is evident from this table that pupils in all types of Faith schools, and in Secular schools that run their own admissions, are at an advantage over pupils in standard non-autonomous Secular

¹⁷ However, relatively few are thought to have done so. West and Hind (2003) provide some evidence on this at secondary level: only about 10% of Voluntary Aided schools were interviewing applicants or their parents.

primary schools, both in terms of initial attainment and background characteristics usually associated with educational disadvantage. Pupils start off in these schools with attainments that are, on average, 1.2 to 1.7 points (1 point is equivalent to one term) ahead of their counterparts in non-autonomous Secular schools. This is around 15% of one standard deviation and about the same as the advantage in terms of final attainment at age-11 reported in Table 2. Certainly, this difference may partly be because these pupils have already spent some time in Faith schools prior to age-7, and may have reaped some educational benefits. However, pupils in Faith and autonomous schools are also much less likely to be on a low income that entitles them to free school meals, more likely to be White and more likely to have English as their first language. The advantage of these schools in terms of lower free school meal entitlement also amounts to 15-20% of one standard deviation, and it is hard to see how such differences in background can be a *consequence* of Faith school attendance.

Some of these disparities can be explained by differences in geographical setting, but not all: Columns 6-10 report the same regressions once we include postcode-level fixed effects, and show that many differences persist even across pupils who live in the same street but attend different types of school. These differences are less marked in terms of ethnicity and languages, but still strong in terms of free school meal entitlement and prior attainment. In Faith schools, pupils still start at the beginning of our period some 0.7 to 1.2 terms ahead of Secular, non-autonomous pupils who live in the same street, and are about 4 percentage points less likely to be eligible for free meals (on a base of 20 percent).

5.3. Regression estimates of progress between ages 7 and 11

Next we turn to regression estimates of the model in Equation (1). Results from our first exercise are shown in Table 4. The dependent variable is the pupil-mean of the Maths and English percentiles described in Table 2. Column (1) provides information on the raw differences between

schools (similar to those in Table 2) by regressing this measure of age-11 attainment on school-type dummies (and academic year dummies).

In Column (2) we control for initial attainment groups at age 7 using dummies for combinations of Levels reached in each of the three subject areas – Maths, Reading and Writing. After taking account of empty cells and aggregating cells with low counts, this gives us 183 dummy variables that classify initial attainment groups. Controlling for age-7 attainment in this way more than halves the differences between mean age-11 attainment of pupils attending different types of school¹⁸. Nevertheless, pupils in Faith schools and autonomous schools still appear to do better, even when starting from the same age-7 base. In Faith, autonomous schools, pupil attainments are nearly 2.5 percentiles above pupils in the same age-7 attainment group in non-selective Secular schools¹⁹.

Column (3) introduces the school and pupil level controls detailed in Appendix Table 7 alongside postcode-of-residence fixed effects to control for unobserved education-related characteristics and preferences that are common to close neighbours. Identification requires multiple school types per postcode, so we restrict the sample to postcodes where this condition holds, leaving us with a much smaller sample. In this case, we are comparing neighbouring pupils (i.e. pupils from the same postcode) with similar characteristics and the same school choice options, but attending different primary school types. As a result, the gap between Secular, non-autonomous schools and other school types is reduced and we now find no evidence of an advantage for pupils in Faith schools over Secular schools when these do not have autonomy over their own admissions. However, pupils emerge with a slightly average higher level of attainment from autonomous

¹⁸ In part this is because the age-7 attainment may in turn be affected by school type, since pupils may spend up to three years in the same school before their age-7 tests.

¹⁹ We have tried other specifications of the value-added model. A common alternative assumption is that $h_{ijt} = (\alpha + \beta_j + x'_{ij}\gamma + \varepsilon_{ij})t$, so that $(h_{ij2} - h_{ij1}) = \alpha + \beta_j + x'_{ij}\gamma + \varepsilon_{ij}$, in which case we can just regress the difference between pupil's age-11 and age-7 point scores on school type dummies and other background characteristics. The results from this exercise convey a similar message to that in Table 4. They are available upon request from the authors.

schools – both Faith and Secular – than they do from schools that are more closely controlled by the LEA. One might suspect that this advantage has still at least in part to do with selection on pupil characteristics that are correlated with progress between ages 7 and 11, some of which we are not able to control for. In fact, the evidence in Table 3 suggests that, even conditional on postcode-of-residence fixed effects, pupils attending different school types are not fully observationally equivalent. However, we cannot rule out the possibility that there are real advantages in the more autonomous governance structures of Voluntary Aided and Foundation schools²⁰.

We argued in Section 4 that secondary school choice (at age-12) is informative about family preferences regarding religious education, as well as pupil adequacy in terms of school admission criteria. Hence, controlling for future school choice could help us to condition out unobservable family-specific factors that jointly determine primary school choice and primary phase educational progress. Bearing this in mind, we estimate the attainment model (1) allowing for secondary-school-type \times postcode-of-residence fixed effects (column (4)) or secondary-school \times postcode-of-residence fixed effects (Column (5)). The results in both specifications reveal a very small residual gap between pupils emerging from autonomous primaries and those from baseline schools – at around 0.8-1 percentile of the pupil distribution or 0.02-0.04 of the standard deviation in test scores. It is worth emphasizing that although this estimate is very modest compared to some of the previous findings in the literature, it could still be upward biased by selection into Faith primary schools amongst pupils who live in the same postcode, and go on to attend the same secondary school²¹.

²⁰ Note that we have re-estimated the specification of Column (2) on the sample of Column (3) to check that the attenuation in the coefficients when we move to from Column (2) to Column (3) is not attributable to a change in composition arising from the reduced sample size. Note also that the attenuation of our estimates is unlikely to be due driven by measurement error and amplification of the signal to noise ratio, since nearly all our variables are dummies based on administrative data. Moreover, the standard errors remain small and there is substantial within-postcode variation in pupils' choice of primary school (the within-postcode variance in school type indicators is 20-30% of the total).

²¹ In fact, the last five columns of Table 3 suggests that, even conditional on secondary school type \times postcode-of-residence fixed effects, pupils attending autonomous primary schools are disproportionately associated with educationally advantageous observable characteristics.

A striking finding emerging from the results so far is that ‘Faith’ affiliation is not, in itself, an indicator of higher educational standards. Faith, autonomous schools have mean attainments that are only 0.15 percentiles higher than Secular, autonomous schools (in Column (5)), and not significantly so (the F-test for equality of the two parameters has a p-value of 0.7236). Moreover, pupils from Faith schools seem to do slightly worse than pupils from Secular schools when admissions are not under their control²². This indicates that religious character is not, on its own, linked to better school performance.

5.4. Regression estimates on samples restricted by future school-sector choice

The evidence provided so far suggests that autonomous primary schools – either Faith or Secular – have a very small effect on pupil achievement at the end of primary schools once we account for unobservable individual and parental heterogeneity by comparing close neighbours attending the same secondary school. We have also pointed out that, despite being very small, estimates of the effect of attending an autonomous primary school conditional on secondary-school \times postcode-of-residence fixed effects might still be upward biased. In this section, we compare estimates for sub-groups of pupils who exhibit very different school choice patterns over both primary and secondary phases. As outlined in Section 4, and borrowing from the intuitions in Altonji et al. (2005), we argue that comparison of the ‘Faith school effect’ for these sub-groups, relative to their degree of sorting on observable characteristics into religious education, is informative about the role that sorting on unobservables plays in generating the small ‘autonomous’ school advantage we have found.

School choice transitions between primary and secondary phases are shown in Appendix Table 8. About 77.5% of pupils in Secular, non-autonomous primary schools transits to Secular,

²² Notice that the null hypothesis that the coefficient for Faith, autonomous schools and that for Secular, autonomous schools are equal to the parameter for Faith, non-autonomous schools are rejected with 0.000 p-value and 0.037 p-value, respectively.

non-autonomous secondary schools. Similarly, 54.8% of those attending a Secular, autonomous primary transits to a Secular, autonomous school for their secondary education. Together, this implies that 54% of our sample stays in the Secular sector in both phases, with just over half (52%) in Community schools controlled by the Local Education Authority. Looking at Faith-school pupils, more than 50% of pupils in Faith, autonomous primary schools (about 10% of our sample) stay within the Faith, autonomous sector during the secondary education phase. On the other hand, some 170,000 pupils (18%) switch out of the Faith sector at the secondary phase, and around 39,000 (4.2%) switch into it. If our conjectures about the relationship between school choice and family background are correct, we would expect to find very different estimates of the Faith-autonomous primary school effect amongst these different groups, and very different patterns of sorting along observable lines into Faith autonomous schools.

We start this investigation by considering the sub-sample of individuals who stay in the same school type during both primary and secondary education (i.e. the pupils in all the diagonal cells of Appendix Table 8). Students from families that commit to religious and autonomous schooling over both educational phases differ widely in terms of their observable (and so presumably unobservable) education-related characteristics from pupils who attend Secular, non-autonomous schools in both phases. We provide supporting evidence in the top panel of Table 5, where we present the results from regressions of individual observable characteristics on school-type dummies for the sub-sample of ‘stayers’. This exercise is similar to the one presented in Table 3. Individuals who stay in Faith and Secular, autonomous schools over both phases of education have higher age-7 achievement than pupils in Secular, non autonomous schools. They are also less likely to be eligible for free school meals and to carry special education needs, and more likely to be of White origin and to speak English as their first language. Importantly, the differences are larger for the sub-sample of Faith-Faith and Secular-Secular ‘stayers’ considered here than for the sample including all pupils (compare Columns 6 to 10 of Table 3). This high degree of observable selection

into Faith primary schooling is indicative that there could also be a large degree of selection on other characteristics that we cannot control for.

Consider now the sample of pupils who change school type on transition from primary to secondary education (i.e. we exclude from our sample pupils in the diagonal cells of the transition matrix of Appendix Table 8). Although we cannot rule out the possibility that pupils making Faith-Secular transitions differ systematically from pupils making Secular-Faith transitions, we expect the differences in background characteristics and school preferences between these groups to be far less marked than the differences between pupils who exhibit Faith-Faith and Secular-Secular school-choice patterns: both groups of ‘switchers’ exhibit similar preferences for Faith schooling, and are (evidently) of a type who is likely to be admitted. Once more, we provide some evidence that this is the case by looking at the differences in observable characteristics in these groups in the bottom panel of Table 5. Although there are still some significant differences between pupils attending Faith-autonomous schools in the primary phase rather than the secondary phase, these differences are much lower in magnitude than for pupils choosing Faith primary schools in the population (Table 3), or choosing Faith-autonomous schools in both phases (Table 5, top panel).

Our argument then is that, if selection on observables provides a guide to selection on unobservables, the bias in our estimates of the effect of Faith-autonomous schooling on test scores should track the degree of selection on observable characteristics in these different sub-groups. Thus, the ‘stayer’ subgroup should yield higher estimates than we found for the whole of the population of pupils, while the ‘switcher’ sub-group should give lower estimates than both for the population of pupils and for the ‘stayer’ sub-sample. We present our empirical findings in Table 6, where we use the same specification as in Table 4, Column (3), on these sub-samples. We focus our discussion on the estimated effects of autonomous schools, since we have already shown that it is only in these schools that there appear to be significant educational advantages.

For the sub-sample of ‘stayers’ attending Faith, autonomous schools in both phases, we find a

test scores advantage of 2.7 percentiles, relative to pupils attending Secular controlled schools in both phases; for Secular, autonomous school pupils the figure is around 2.2 percentiles. These estimates are 60-100% higher than those found for the full sample in Table 4. On the other hand, we now find that the performance gap of Faith and autonomous schools for the sub-sample of pupils who ‘switch’ school-type at the end of primary school is below zero, although not significant. Pupils who attended an autonomous primary up to age 11, but move to a non-autonomous secondary school, perform no better – if not marginally worse – at age 11 than pupils who *do not* attend an autonomous primary school up to age-11, but go on to attend an autonomous secondary thereafter. These findings seem to count against Faith and autonomous schools being the driving factor behind better test results: the estimated advantage of Faith and autonomous school attendance closely tracks the degree of selection on observable characteristics, which makes pupil sorting a much more plausible explanation for any difference in average pupil performance²³.

We can further see this by considering a few simple descriptive statistics. Looking into Table 3 to Table 6, it can be seen that pupils in Faith-autonomous schools are about 10-15% of one standard deviation above pupils in Secular controlled schools in terms of observable characteristics that are associated to higher age-11 test scores (e.g. age 7 test scores, not on free meals, no special educational needs). In turn, they have a value-added advantage of around 6% of one standard deviation, about half the advantage in terms of predetermined observable characteristics. For the sub-sample of ‘stayers’, the advantage in terms of observable characteristics is about 20% of one standard deviation, while the advantage in value-added about 10%. On the other hand, for the sample of ‘switchers’, the *disadvantage* in age-7 attainments and free-meal entitlement (the most important drivers of age-11 test scores in our models) is about 2.5% of one-standard deviation, while the *disadvantage* in value-added, although insignificant, is about 1% of one-standard

²³ Incidentally, our evidence does not point to any beneficial impact of attending a Faith primary school for more disadvantaged pupils, like those with special education need status. We also experimented breaking down the sample by free school meal eligibility (usually associated with economically disadvantaged family background), but still failed to find any differential impact. This is at odds with most of the US based evidence.

deviation. There is clearly a close relationship between selection on observables into Faith autonomous schools and the magnitude of our estimates of their performance advantage, which suggests that most of the documented effects of Faith and autonomous schooling are likely to be driven by school-side selection and pupil sorting.

To conclude, we formalize our arguments using a variant of the Altonji et al. (2005) procedure outlined in Section 4. We estimate the model set out in Equation (2) on the different sub-samples of ‘stayers’ and ‘switchers’ with varying constraints on ρ . We use specifications in which we *exclude* postcode fixed-effects (and school characteristics) from our regressions, but include pupil age-7 test scores and background characteristics²⁴. We also consider only pupils choosing Faith-autonomous or Secular-controlled schools, since we need a dichotomous ‘treatment’ variable. This approach allows us to estimate the degree of selection on observables into Faith autonomous schools for ‘stayers’ and ‘switchers’ (i.e. the correlation between $x'_{ij}\hat{\gamma} + \hat{g}_1(h_{ij1})$ and $x'_{ij}\hat{\lambda} + \hat{g}_2(h_{ij1})$ described in Section 4), and calibrate the degree of selection on unobservables through our choice of ρ .

Using this method, we find that setting ρ to 0.09 in regressions that consider the whole of the population of pupils drives the estimates of the Faith-autonomous school advantage to zero (compare with the coefficients in Column (2) of Table 4). For the sub-sample of ‘stayers’, we need ρ to be 0.11 in order to drive the Faith, autonomous school effect to zero. Finally, without postcode fixed effects, estimates on our sub-sample of ‘switchers’ yields an estimated Faith-autonomous effect of about 1 percentile (rather than the marginally negative coefficient of Column (2) of Table 6), but this too is driven to zero if ρ is set to just 0.04. In comparison, the correlation between the predictions from the test score and selection equations (i.e. the degree of selection on observables) is 0.42 for the population of pupils, 0.46 in the group of ‘stayers’ and 0.04 in the group of ‘switchers’. Clearly, for the population of pupils and subsample of ‘stayers’, we can explain most

²⁴ We use a modified version of the two-step `treatreg` command in Stata 10. We need to exclude the postcode fixed effects, because the model requires a probit first stage.

of the positive and significant coefficients of Table 4 and Table 6 with only a very small amount of selection on unobservables, relative to the selection on observables that we can observe in our data. Moreover, our group of ‘switchers’ exhibits a very low degree of selection on observables, even without controls for place of residence. A comparable amount of selection on unobservables would imply no causal Faith-autonomous school effect for these pupils²⁵.

Note that, in the analysis so far, we have assumed that secondary schools do not screen pupils on the basis of *progress* during primary school. As we stated in Section 4, this assumption is in line with previous evidence, and plausible in our context because we have dropped from our analysis all ‘Grammar’ schools that have admissions tests, and because pupils in England secure their secondary school places before sitting for their final primary-phase assessments²⁶. Suppose however that Faith secondary schools were able to select pupils on the basis of information about educational progress that they can observe, while we cannot. This school-side selection could mean that pupils who switch from Faith primary schools to secular secondary schools are excluded from Faith secondary schooling because they progressed too slowly during the primary phase. Conversely, pupils who switch from secular primary schools to Faith secondary schools might be those that progressed unexpectedly rapidly during primary education, and so are picked out of the secular sector by Faith secondary schools. In this case, selection on *unobservable* value-added generating attributes would have the opposite sign to selection on the basis of *observable* value-added generating characteristics within our ‘switchers’ sub-sample: pupils with the most educationally advantageous observable characteristics switch out of the Faith sector at the end of primary education, whilst pupils with the most educationally advantageous unobservable (to us)

²⁵ Obviously, if we forced ρ to assume negative values, i.e. if we assumed negative selection on unobservables, then we would find estimates that lie above those documented in Table 4 and Table 6. However, negative sorting is very unlikely in England given the structure of the school system and the evidence provided above.

²⁶ Grammar schools educate around 6-7% of secondary school pupils in England and are predominantly non-religiously affiliated. There are no Grammar primary schools.

characteristics switch in. Clearly then, the estimates based on our switchers sub-sample will be downward biased.

Similarly, we have treated the benefits of Faith schooling as homogenous across pupils. This is explicit in equations (1) and (2) (and implicit in previous studies that use instrumental variables approaches). Suppose instead that returns to Faith schooling are heterogeneous and pupils choose school types according to the expected returns. In this case, the lowest estimates based on our ‘switchers’ sample may be downward biased, because pupils benefiting least may switch out of Faith education. On the other hand, our highest estimates based on those who stay in Faith education – presumably those who benefit from it the most – are upward biased.

Nevertheless, in either case – school-side selection on the basis of unobservable value-added generating characteristics or heterogeneous returns – estimates based on the samples of ‘switchers’ and ‘stayers’ still allow us to provide a lower and an upper bound respectively for the likely effect of Faith schooling on educational achievement in the absence of school selection and pupil sorting.

5.5. Different religious affiliation

So far, we have brushed aside consideration of any denominational differences between schools. However, denominational differences might be an important consideration, and there are reasons why we may expect to see differences in educational outcomes between schools of different Faiths. For example, the ethos of one Faith may have more educational impact than another, or non-random selection of pupils into schools may be greater within one denomination than another²⁷. Almost all Faith schools in England are either Church of England or Catholic, so we can say little about minority Faiths (Jewish, Muslim etc.). However, given the data at hand, we were able to compare

²⁷ E.g., Noden et al. (1998) found that Roman Catholic families express a strong preference for Roman Catholic schools.

Catholic and Church of England schools²⁸. This is also interesting in reference to the US literature, which focuses on Catholic (private) schools.

In a nutshell, we found some evidence that the attainment gap between the Faith and Secular sector schools is attributable to pupils in Catholic schools in the Faith, autonomous sector. However, even then, most of this difference can be attributed to unobservable differences between Catholic and non-Catholic pupils. The gap only occurs between pupils who choose to attend a Catholic school at *any* stage in their school careers, and those who never do so. Nothing from our analysis of the differential impacts of Catholic and Church of England schools leads us to conclude that our main estimates are misleading about the overall impact of attending a Faith school.

5.6. Instrumental variable estimates

We began our paper casting doubts on the validity of those IV approaches that instrument school choice using characteristics of residential neighbourhoods, such as availability of religious schools or fraction of individuals belonging to some religious group. We argued that this is a bad strategy because families choose where to live – often for reasons related to access to good schools – so that neighbourhood characteristics are likely to be correlated with family characteristics, preferences and pupil achievements. In fact, instrumenting school choice with characteristics that are the result of the sorting equilibrium arising from school choice is likely to inflate any bias rather than fix it (for reasons similar to those discussed in Altonji et al., 2002, 2005).

To evaluate this approach, we (partly) replicated Neal (1997), and instrumented religious school attendance using either the supply of religious schools (relative to Secular schools) in the neighbourhood of residence, or the fraction of Christians in the neighbourhood population. Although we find that the instruments are powerful in the first stage equation, the second stage estimates of the effect of attending a Faith, autonomous school are too large to be credible estimates

²⁸ Catholic schools account for about 53% of the Faith, autonomous sector in England, but there are almost none in the non-autonomous sector.

of the causal impact of attending a religious school, given the range of mean school performance in the data²⁹. More importantly, the IV coefficients are always well above the basic OLS estimates without any controls presented in Table 4. To believe these estimates, one would have to assume that there is strong negative selection into Faith schools based on unobservables, and that pupils with educational disadvantageous characteristics choose Faith schooling (in line with the claims about Catholic schools in the US literature, e.g. Neal (1997)). This negative sorting is very unlikely in England given what we know about the school system, and the evidence on the way pupils are sorted into Faith schools on the basis of observable characteristics provided in Section 5. At least in the English setting, these instruments are invalid and simply magnify the effect of unobservable parental preferences and school-side selection.

6. Conclusions

We have provided a number of estimates of the effect of attending a Faith school in England on pupils' educational progress between ages 7 and 11. Our approach has deliberately avoided instrumental variable strategies adopted by previous work in the field, because we do not believe (at least for the English setting) that there are any credible instruments for Faith school attendance that are uncorrelated with family background, either directly or through residential sorting. Instead we have exploited the fact that we have around one million pupils in our database, which, in conjunction with precise details about place of residence, academic record and future school choice, allows us to control quite carefully for factors that influence the propensity to attend Faith schools.

We make no claim to have put a precise number on the causal impact of Faith school attendance, and have demonstrated that the magnitude of any difference between Faith school pupils and Secular school pupils depends substantially on the way we cut the sample. What then are we to make of these results? One thing that seems clear is that, although Faith schools – and other

²⁹ These were around 5-6 percentiles when using the relative supply of religious education in the neighbourhood as an instrument, and up to 17 percentile when using the fraction of Christian in the local population.

schools that have ‘autonomous’ admissions and governance arrangements – tend to admit ‘better’ pupils, there is no unambiguous performance advantage that cannot be attributed purely to pupil-side sorting into these schools, or to school-side selection of pupils likely to show the fastest progress. Pupils who attended Faith or autonomous schools at primary phase, but not at the secondary phase, do no better in primary school than pupils who attend Faith or autonomous schools at the secondary phase, but not at the primary phase. The Faith or autonomous-school gap in attainments at primary phase seems largely attributable to differences between those pupils who choose to attend such schools at *any* stage in their educational careers, and those who choose never to do so or are excluded from doing so by school selection procedures.

In any case, we find no evidence that Faith affiliation lies behind the test-score advantage commonly attributed to Faith schools in England. A generous reading of the results suggests that pupils in schools that have more autonomous governance and admissions structures – a set that includes Faith schools – do progress marginally faster. A pupil starting in an autonomous school at age 7 could expect to be one percentile higher in the distribution of pupil attainments by age-11 than a comparable pupil attending a standard Secular-controlled school, even when these two pupils live in the same postcode and go on to choose the same secondary school. Our upper bound estimates put this figure at 2.7 percentiles. To put this in perspective, we draw on results in Machin and McNally (2004) that report labour market returns to age-10 reading tests, based on the British Cohort Survey³⁰. Their figures (reported in their Table 7) indicate that the labour market return to a one percentile move up the attainment distribution at age 10 was around 0.0042%, conditional on family background. In other words, the labour market impact of these small school quality differences seems very slight. Certainly, the cumulative effect over 12 years of compulsory schooling could be more substantial than this would suggest, and there may be other impacts from schooling of religious ethos – on staying on rates and child wellbeing for example – that are outside

³⁰ The British Cohort Survey follows a cohort of children born in one week in 1980 through to adulthood. The reading tests were administered in 1990 when the children were aged 10.

the scope of this study. However, pupils in Faith schools that are under close Local Education Authority control do not progress any faster than similar pupils in comparable Secular schools. Any performance impact from 'Faith' schools in England seems to be closely linked to their autonomous governance and admissions arrangements, and not to religious character.

7. References

- Allen, R. (2006), "Allocating pupils to their nearest secondary school: The consequences for social and ability stratification", Bedford Group for Lifecourse and Statistical Studies Occasional Paper No.3.
- Altonji, J., T. Elder and C. Taber (2002), "An evaluation of instrumental variable strategies for estimating the effects of Catholic schools", National Bureau of Economic Research (NBER) Working Paper 9358.
- Altonji, G., T. Elder and C. Taber (2005), "Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools", *Journal of Political Economy*, 113(1):151-184.
- Barro, R and R. McCleary (2003), "Religion and economic growth", NBER Working Paper 9682,
- Benton, T., D. Hutchinson, I. Schagen and E. Scott (2003), "Study of the performance of maintained secondary schools in England", National Foundation for Educational Research (NFER) working paper.
- Black, S. (1999), "Do better schools matter? Parental valuation of elementary education", *Quarterly Journal of Economics*, 114, 578-599.
- Coleman, J., T. Hoffer and S. Kilgore (1982), *High school achievement: Public, Catholic and private schools compared*, New York; Basic Books.
- Dronkers, J. (2004), "Do public and religious schools really differ? Assessing the European evidence", in: Wolf, P.J and Macedo, S. (Eds.), *Educating Citizens. International Perspectives on Civic Values and School Choice*, Brookings Institution Press, Washington, DC.
- Evans, W. and R. Schwab (1995), "Finishing high school and starting college: Do Catholic schools make a difference?", *Quarterly Journal of Economics*, 110, 941-974.
- Figlio, D. and J. Stone (1999), "School choice and student performance: Are private schools really better?", *Research on Labour Economics*, 18, 115-140.
- Gibbons, S. and S. Machin (2006), "Paying for primary schools: Supply constraints, popularity or congestion?", *Economic Journal*, 116, C77-C92.
- Grogger, J. (1995), "The effect of arrests on the employment and earnings of young men", *Quarterly Journal of Economics*, 110, 51-71.
- Grogger, J. and D. Neal (2000), "Further evidence on the effects of Catholic secondary schooling", *Brookings-Wharton Papers on Urban Affairs*, 151-201.
- Gruber, J. (2005), "Religious market structure, religious participation and outcomes: Is religion good for you?", *Advances in Economic Analysis and Policy*, 5-30.
- Guiso, L., P. Sapienza and L. Zingales (2003), "People's opium? Religion and economic attitudes", *Journal of Monetary Economics*, 50, 225-282.
- Hanushek, E. (2003), "The Failure of input-based school policies", *Economic Journal*, 113, F64-98.
- Heckman, J. (2000), "Policies to foster human capital", *Research in Economics*, 54, 3-56.

- Jepsen, C. (2003), "The effectiveness of Catholic primary schooling", *Journal of Human Resources*, 38, 928-941.
- Hoxby, C. (2004), "School choice and school competition: Evidence from the United States", *Swedish Economic Policy Review*, 10(2).
- Kane, T., D. Staiger and S. Reigg (2005), "School quality, neighbourhoods and housing prices: The impacts of school desegregation", NBER Working Paper 11347.
- Le Grand, J. (1991), *Equity and choice*. London, Harper Collins.
- Le Grand, J. (1993), *Quasi-markets and social policy*. London, Macmillan.
- Machin, S. and A. Vignoles (2005), *What's the Good of Education?*, Princeton University Press.
- Machin, S. and S. McNally (2004), "The literacy hour", Centre for Economics of Education (CEE) Discussion Paper 43.
- Manning, A. and S. Pischke (2006), "Comprehensive versus selective schooling in England and Wales: What do we know?", CEE Discussion Paper 66.
- Neal, D. (1997), "The effects of Catholic secondary schooling on educational achievement", *Journal of Labour Economics*, 15, 98-123.
- Noden, P., A. West, M. David and A. Edge (1998), "Choices and destinations at transfer to secondary schools in London", *Journal of Education Policy*, 13(2), 221-236
- Noell, J. (1982), "Public and private schools: A reanalysis of 'Public and private schools'", *Sociology of Education*, 56, 12-132.
- Prais, S., (2005), "The superior educational attainments of pupils in religious foundation schools in England", *National Institute Economic Review*, 193, 102-105.
- Sander, W. (1996), "Catholic grade schools and academic achievement", *Journal of Human Resources*, 31, 540-548.
- Schagen, S., D. Davies, P. Rudd and I. Schagen (2002), "The impact of specialist and Faith schools on performance", LGA Educational Research Programme Report 28, Slough, NFER.
- Tiebout, C. (1956), 'A pure theory of local expenditures', *Journal of Political Economy*, 64(5).
- Weber, M. (1905), *The Protestant Ethic and the Spirit of Capitalism* (2003, reproduction of 1956 translation of 1905 original), Dover Publications.
- West, A. (2005), "School choice: The limits of quasi regulation", *Consumer Policy Review*, 15, 94-98.
- West, A. and A. Hind (2003), "Secondary school admissions in England: Exploring the extent of overt and covert selection", Report for Research and Information on State Education Trust, www.risetrust.org.uk/admissions.html

Table 1: Institutional characteristics of primary schools in England

Type	Faith	Governors (approximately)	Admissions authority	Assets owned by	Employer
Community	Secular	Parents >30%, Staff <30%, LEA 20%, Community 20%	LEA	LEA	LEA
Foundation	Mostly Secular, some C. of E.,	Parents >30%, Staff <30%, Foundation/Partnership <25%, LEA <20%, Community 10%	Governors	Foundation or Governors	Governors
Voluntary Aided	Mostly C. of E. or Catholic, some other Faith, some Secular	Foundation >50%, Parents >30%, LEA <10%, Staff <30%	Governors	Foundation	Governors
Voluntary Controlled	Mostly C. of E., some other Faith, some Secular	Parents >30%, Staff <30%, Foundation <25%, LEA <20%, Community 10%	LEA	LEA	LEA

Note: C. of E. means Church of England.

Table 2: Age-11 attainments by school type; descriptive statistics

Variable	Mean	Std.Dev.	Percentage of age-11 pupils
Average KS2 score, Mathematics and English (percentiles)	50.50	26.61	100%
Faith (non-autonomous or autonomous)	53.85	26.17	29.21%
Faith, non-autonomous	52.43	26.46	9.94%
Faith, autonomous	54.58	25.99	19.27%
Secular (non-autonomous or autonomous)	49.12	26.66	70.79%
Secular, non-autonomous	49.00	26.67	68.18%
Secular, autonomous	52.17	26.78	2.61%
Autonomous (Faith or Secular)	54.30	26.03	21.88%

Note: the total number of observations is 929,955. Pupils attending or moving to schools with other religious denominations are dropped from the sample; they amount to about 0.6% of the sample (6,387) pupils. Autonomous schools include (Secular and Faith) Foundation and Voluntary Aided schools. Non-autonomous schools include Community and Voluntary Controlled schools.

Table 3: Prior attainment and pupil background by primary school type

	<i>No controls</i>					<i>Postcode fixed effects</i>					<i>Postcode × Secondary school fixed effects</i>				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
	Age-7 KS1 points	Free meal eligible	SEN status	White	English first language	Age-7 KS1 points	Free Meal Eligible	SEN status	White	English first language	Age-7 KS1 points	Free meal eligible	SEN status	White	English first language
Mean (Std.Dev.)	44.752 (10.212)	0.163 (0.369)	0.208 (0.406)	0.845 (0.361)	0.903 (0.296)	44.695 (10.198)	0.181 (0.385)	0.218 (0.413)	0.860 (0.346)	0.924 (0.265)	44.894 (10.135)	0.149 (0.356)	0.212 (0.409)	0.902 (0.297)	0.950 (0.217)
Faith, autonomous	1.657 (0.075)	-0.046 (0.003)	-0.014 (0.002)	0.033 (0.005)	0.037 (0.004)	1.283 (0.053)	-0.041 (0.001)	-0.013 (0.002)	0.024 (0.002)	0.034 (0.002)	0.231 (0.076)	-0.016 (0.002)	0.006 (0.003)	0.010 (0.002)	0.016 (0.002)
Faith, non- autonomous	1.290 (0.097)	-0.080 (0.003)	-0.014 (0.003)	0.074 (0.006)	0.061 (0.005)	0.133 (0.078)	-0.014 (0.002)	0.004 (0.003)	0.005 (0.002)	0.009 (0.002)	-0.089 (0.093)	-0.007 (0.002)	0.010 (0.004)	0.007 (0.003)	0.006 (0.002)
Secular, autonomous	1.209 (0.232)	-0.067 (0.007)	-0.022 (0.006)	0.036 (0.014)	0.029 (0.012)	0.691 (0.150)	-0.025 (0.004)	-0.024 (0.005)	0.012 (0.005)	0.005 (0.006)	0.311 (0.191)	-0.018 (0.005)	-0.004 (0.007)	0.014 (0.005)	-0.002 (0.006)
Secular, non- autonomous	44.273 (0.042)	0.181 (0.002)	0.213 (0.001)	0.831 (0.003)	0.889 (0.003)	44.254 (0.026)	0.196 (0.001)	0.222 (0.001)	0.852 (0.001)	0.912 (0.001)	44.833 (0.039)	0.156 (0.001)	0.209 (0.001)	0.898 (0.001)	0.945 (0.001)

Note: The top part of the table shows raw means and standard deviations for all schools. The bottom part shows means for Secular, non-autonomous schools, and mean differences for other school categories with respect to Secular, non-autonomous schools. Means and mean differences in the bottom part of the Table are obtained from regressions at the pupil level without controls or controlling for postcode fixed effects; standard errors clustered at the schools level. SEN means: Special Educational Needs (with and without statements). Sample size: no controls 929,955; Postcode fixed effects: 281,408; Postcode × Secondary school fixed effects: 101,199.

Table 4: School type and mean age-11 attainment;
conditional on initial attainment, background and place of residence

	(1)	(2)	(3)	(4)	(5)
Faith, autonomous	5.582 (0.212)	2.338 (0.158)	1.686 (0.126)	0.962 (0.148)	0.817 (0.168)
Faith, non-autonomous	3.425 (0.274)	0.918 (0.199)	0.023 (0.166)	-0.147 (0.174)	-0.222 (0.193)
Secular, autonomous	3.168 (0.606)	0.925 (0.453)	1.118 (0.337)	0.973 (0.371)	0.671 (0.397)
Age-7 attainment	No	Yes	Yes	Yes	Yes
Individual and school level controls	No	No	Yes	Yes	Yes
Postcode fixed effects	No	No	Yes	No	No
Postcode × Secondary school type fixed effects	No	No	No	Yes	No
Postcode × Secondary school fixed effects	No	No	No	No	Yes
Schools	14,821	14,821	14,020	13,357	12,089
Observations	929,955	929,955	281,408	155,085	100,199

Note: Regressions at the pupil level; standard errors clustered at the primary school level. Baseline: Secular, non-autonomous schools. Controls with descriptive statistics are listed in Appendix Table 7.

Table 5: Prior attainment and pupil background by primary school type; pupils who stay or switch school types across primary and secondary phases

	(1)	(2)	(3)	(4)	(5)
	Age-7 KS1 points	Free meal eligible	SEN status	White	English first language
<i>Panel A: Stayers</i>					
Faith, autonomous	2.233 (0.071)	-0.062 (0.002)	-0.034 (0.003)	0.031 (0.003)	0.037 (0.002)
Faith, non-autonomous	-1.007 (0.730)	0.015 (0.021)	0.006 (0.028)	0.035 (0.017)	0.002 (0.016)
Secular, autonomous	1.826 (0.292)	-0.039 (0.010)	-0.091 (0.012)	0.010 (0.011)	0.002 (0.011)
Secular, non-autonomous	43.90 (0.028)	0.216 (0.010)	0.228 (0.001)	0.850 (0.001)	0.908 (0.001)
<i>Panel B: Switchers</i>					
Faith, autonomous	-0.268 (0.111)	-0.019 (0.003)	0.014 (0.004)	0.008 (0.004)	0.026 (0.003)
Faith, non-autonomous	-0.576 (0.131)	-0.011 (0.004)	0.015 (0.005)	0.011 (0.004)	0.016 (0.003)
Secular, autonomous	0.171 (0.264)	-0.012 (0.008)	-0.000 (0.011)	0.006 (0.009)	0.017 (0.009)
Secular, non-autonomous	45.07 (0.069)	0.166 (0.002)	0.213 (0.003)	0.856 (0.003)	0.917 (0.002)

Note: Table shows means for Secular, non-autonomous schools, and mean differences for other school categories with respect to Secular, non-autonomous schools. Means and mean differences are obtained from regressions at the pupil level with postcode fixed effects; standard errors clustered at the schools level. SEN means: Special Educational Needs (with and without statements). Panel A only includes pupils who attend the same type of schools in both periods. Panel B excludes pupils who attend the same types of school in both phases.

Table 6: School type and mean age-11 attainment; ‘stayer’ and ‘switcher’ sub-samples

	(1)	(2)
	Including stayers across both phases only	Excluding stayers across both phases within all school types
Faith, autonomous	2.672 (0.183)	-0.106 (0.225)
Faith, non-autonomous	2.036 (2.050)	-1.319 (0.248)
Secular, autonomous	2.176 (0.594)	-1.293 (0.513)
Age-7 attainment	Yes	Yes
Individual and school level controls	Yes	Yes
Postcode fixed effects	Yes	Yes
Schools	10,535	9,956
Observations	170,931	110,481

Note: Regressions at the pupil level; standard errors clustered at the primary school level. Baseline: Secular, non-autonomous schools. Column 1 only includes pupils who attend the same type of schools in both periods. Column 2 excludes pupils who attend the same types of school in both phases. Controls with descriptive statistics are listed in Appendix Table 7.

8. Appendix Tables

Table 7: Control variables: descriptive statistics

Variable	Mean	Std.Dev.	Min,Max
<i>Pupil Level</i>			
Female	0.496	0.499	0,1
Native language English	0.902	0.296	0,1
Native language not available	0.022	0.148	0,1
Native language not English	0.075	0.263	0,1
Pupil eligible for free school meals (FSM)	0.163	0.369	0,1
FSM eligibility status missing	0.022	0.146	0,1
Pupil with special education needs (SEN)	0.208	0.406	0,1
SEN status missing	0.022	0.147	0,1
White ethnicity	0.845	0.361	0,1
Black Caribbean ethnicity	0.014	0.116	0,1
Black Other ethnicity	0.016	0.124	0,1
Indian ethnicity	0.019	0.136	0,1
Pakistani ethnicity	0.023	0.149	0,1
Other Asian ethnicity	0.011	0.103	0,1
Other and mixed ethnicities	0.027	0.161	0,1
Missing ethnicity	0.046	0.211	0,1
Academic Year 2001/2002	0.499	0.500	0,1
<i>School Level</i>			
Total number of pupils	315.8	132.5	13,1292
Pupil/teacher ratio	23.14	3.096	4.3,72.2
Fraction of pupils eligible for FSM	0.169	0.145	0,0.94
Fraction of pupils with SEN	0.197	0.095	0,0.79
Fraction of Whites in school	0.844	0.254	0,1
Fraction of Caribbean Blacks in school	0.013	0.047	0,0.79
Fraction of Other Blacks in school	0.016	0.053	0,1
Fraction of Indians in school	0.019	0.070	0,1
Fraction of Pakistani in school	0.023	0.095	0,1
Fraction of Other Asian in school	0.011	0.053	0,1
Fraction of other and mixed ethnicity in school	0.027	0.051	0,1
Fraction with missing ethnicity in school	0.047	0.167	0,1
Ratio of ethnically classified to total pupils in school	0.409	0.431	0,1

Table 8: Transition matrix between primary and secondary phase, by school type

Current school (age 11)	Future school (age 12)			
	Faith, autonomous	Faith, non-autonomous	Secular, autonomous	Secular, non-autonomous
Faith, autonomous	51.2 (91,774)	0.8 (1,526)	10.8 (19,408)	37.2 (66,497)
Faith, non-autonomous	6.9 (6,343)	2.2 (2,044)	18.5 (17,147)	72.4 (66,948)
Secular, autonomous	5.4 (1,310)	0.5 (122)	54.8 (13,295)	39.3 (9,531)
Secular, non-autonomous	5.2 (32,714)	0.8 (4,785)	16.5 (104,897)	77.5 (491,617)

Note: The table presents cell percentages; total numbers in parentheses.