Economics of Education Review 66 (2018) 154–166

Contents lists available at ScienceDirect

Economics of Education Review

journal homepage: www.elsevier.com/locate/econedurev

Achievement effects from new peers: Who matters to whom?

Duncan McVicar[a,c,*], Julie Moschion[b,c], Chris Ryan[b]

[a] Queen's Management School, Queen's University Belfast, Riddel Hall, 185 Stranmillis Road, Belfast BT9 5EE, United Kingdom
[b] Melbourne Institute of Applied Economic and Social Research, University of Melbourne, Vic 3010, Australia
[c] IZA, Bonn

ARTICLE INFO

Keywords:
Peer effects
School achievement
Education

JEL: I21
I24
J24

ABSTRACT

This paper presents estimates of achievement-related peer effects on school students' literacy using data from national test scores, across multiple literacy measures and student cohorts, for the population of public secondary school students in Years 7 and 9 (aged 12/13 and 14/15 years) in the Australian state of Victoria. Identification is achieved via individual fixed effects and by distinguishing between secondary school peers who attended the same primary school as the individual and those who did not. Estimates of peer effects are based on the new peers, whose primary school achievement could not have been affected by the individual. The results provide strong evidence for the existence of peer effects, with small but positive and statistically significant effects from having higher-achieving peers on average and from having a higher proportion of very high-achieving peers. Further, it is individuals in the middle of the ability distribution who benefit most from having high achieving peers.

1. Introduction

Peer effects refer to externalities in which the actions or characteristics of a reference group affect an individual's behaviour or outcomes. They have been studied in numerous contexts. This paper examines a specific form of peer effects, related to the effect on a student's achievement of the achievement of his or her peers, which if the effects are of sufficient magnitude, have critical implications for students, parents, schools and policy makers. Effects from peer averages imply that parents can improve their child's expected educational achievement by selecting a school with a higher-ability intake, and reforms introducing greater school choice could widen educational inequalities. Non-linear peer effects have additional efficiency implications, such that schools may be able to improve the average achievement of their students, and policy makers the efficiency of the schooling system as a whole, by manipulating the allocation of students across classes or schools. As a result, peer effects in school achievement have attracted an enormous amount of attention in the literature. Establishing the existence and magnitude of peer effects, however, is beset by practical difficulties (Angrist, 2014; Manski, 1993, 2000; Moffitt, 2001), a fact that has probably also contributed to the longevity and ubiquity of the literature. Despite the vast number of studies, disagreement as to the nature, magnitude and even the existence of peer effects in school achievement remains.

The two key difficulties in estimating achievement peer effects arise from endogenous sorting into or within schools and the fact that the individual is a peer of their peers and may therefore influence their peers – the reflection problem. One branch of the school achievement peer effects literature has attempted to address the second challenge by utilising measures of peer quality that pre-date any potential social interactions between the individual and their peers. A promising new strand of this literature exploits the transition between primary and secondary school (so far only in England), and the fact that most of the secondary school peers of any individual attended a different primary school (Gibbons & Telhaj, 2016; Lavy, Silva, & Weinhardt, 2012; Mendolia, Paloyo, & Walker, 2018; Zhang, 2016). Specifically, these papers use the prior performance of peers when they were in different schools – a measure that is immune from reflection problems – to measure the impact of peers' earlier achievement on current student performance. Lavy et al. (2012) also exploits multiple measures of individual student achievement in various learning domains to remove any fixed student-specific effect on achievement, which mitigates confounding effects from sorting into schools, making for a very strong identification strategy. They find no evidence of an average peer effect on individual achievement. Gibbons and Telhaj (2016), Zhang (2016) and Mendolia et al. (2018), however, all find some evidence of small, positive average peer effects. All four studies use data drawn from the same underlying data set; the English National Pupil Database.

* Corresponding author.
E-mail addresses: d.mcvicar@qub.ac.uk (D. McVicar), moschion@unimelb.edu.au (J. Moschion), ryan.c@unimelb.edu.au (C. Ryan).

https://doi.org/10.1016/j.econedurev.2018.07.008
Received 28 September 2017; Received in revised form 18 July 2018; Accepted 31 July 2018
Available online 01 August 2018

The research design adopted in this paper also exploits this transition, when students encounter new classmates at the start of secondary school who did not attend the same primary school. The secondary school achievement of individuals is regressed on the primary school achievement of their secondary school new peers, for whom there is no reflection problem. The approach also makes use of multiple measures of individual achievement, specifically a set of four language or literacy-related measures of achievement. This enables use of an individual fixed effects approach to account for common individual effects on achievement across the four language-related subjects of reading, writing, grammar and spelling. Because the data cover multiple cohorts of students, school-subject fixed effects can also be included to wash out time-invariant correlated effects that might differ across subjects. In other words, peer effects are estimated by exploiting within-individual variation in the performance of new peers across the literacy domain when they were in primary school.

In contrast to the other papers in this strand of the literature, however, this paper uses data from Australian schools rather than English schools and focuses on outcomes at age 12/13 rather than outcomes at age 13/14 or later. In doing so the paper makes a standalone contribution to the knowledge base supporting education policy makers by shedding light on how well the conclusions of these earlier English studies generalise to a similar but not identical educational context. Its single most important contribution, however, is to present evidence of peer effects, including small but non-trivial average peer effects, using a strong identification approach exploiting within-individual variation in test scores. A new placebo test is also added.

Specifically, this paper estimates peer effects using administrative data on the test scores of students attending public schools in the Australian state of Victoria. The data come from the National Assessment Program – Literacy and Numeracy (NAPLAN) conducted across Australia, which provides test scores for five subjects – numeracy, reading, spelling, grammar, and writing – for all students in grades 3, 5, 7, and 9. The national testing system which generates these data was introduced across Australia in 2008, with tests taking place each year in the specific grades, and the paper exploits data from 2008 to 2013. This provides four cohorts of Year 5 (primary school) students subsequently observed in Year 7 (in secondary school), and two cohorts with students also observed in Year 9.

In addition to estimating average peer effects, this paper also assesses the impact of having a high proportion of peers who were in either the top or bottom 10% of the achievement distribution, while also allowing these effects to differ between genders and across the distribution of individual ability (that is, allowing for non-linear and heterogeneous peer effects). This allows testing of empirical support in Australian secondary schools for a wide range of models including the 'bad apple' (disruptive students harm everyone), the 'shining light' (excellent students provide a great example for all), the 'invidious comparison' (outcomes are harmed by the presence of better achieving peers), and ultimately to assess the possible benefits of tracking, at least at the school grade level (see Sacerdote, 2011).

The remainder of this paper is set out as follows. The next section briefly reviews those studies in the school achievement peer effects literature to which this paper most closely speaks. The data are described in more detail in Section 3, along with further discussion of the approach to estimation. The main results are discussed in Section 4, with extensions in Section 5 and concluding discussion in Section 6.

2. Related literature

Recent studies on peer effects in school achievement have looked for both average peer effects and peer effects from and to different points in the ability distribution, using a variety of strategies to overcome the identification problems associated with reflection and endogenous sorting. One strand of this literature exploits random or quasi-random assignment to new peer groups stemming from programs like Metco in the US (Angrist & Lang, 2004) or the Extra Teacher Program in Kenya (Duflo, Dupas, & Kremer, 2011). These studies make credible claims to estimating causal peer effects (or more precisely the absence of causal peer effects) that are internally valid. But a dearth of suitable similar experiments in other contexts means we cannot yet draw general conclusions from this strand of literature alone. A second (non-experimental) group of studies regresses own achievement on lagged peer achievement under the assumption that an individual's current achievement cannot impact her peers' prior achievement (e.g. Atkinson, Burgess, Gregg, Propper, & Proud, 2008; Hanushek, Kain, Markman, & Rivkin, 2003; Lefgren, 2004; Vigdor & Nechyba, 2007). Manski (1993) questions the extent to which this approach truly solves the reflection problem on the grounds that peer scores are likely to be persistent. Additional identification steps – typically school fixed effects and individual controls – are also required to try to wash out sorting effects, arguably with mixed success. Angrist (2014) also critiques this approach for being susceptible to 'mechanical' biases due, among other things, to measurement error and negative intra-group correlation between own test score and the leave-out-mean test score of peers. A third group of studies exploits demographic variation in peer group composition to instrument for peer achievement (e.g. Goux & Maurin, 2007), with the extent to which the resulting estimates can be interpreted as causal peer effects hanging on the plausibility of the exclusion restrictions. Sacerdote (2011) provides a review, summing up the body of evidence from this literature as "mixed" .

The current paper contributes to a fourth group of studies – Gibbons and Telhaj (2016), Lavy et al. (2012), Zhang (2016) and Mendolia et al. (2018) – which arguably falls somewhere between the first and second literature strands. Like the second group of studies described above, these studies also regress own achievement on lagged (school grade) peers' achievement, with a number of additional identification steps to address sorting and other correlated effects. Like the first set of studies described above, however, they focus on the impact of being grouped with *new* peers, in this case as a result of the transition between primary and secondary school in England.[1] These studies are immune to reflection problems because the prior scores (in primary school) of new peers (in secondary school) cannot have been impacted by the individual. This research design also mitigates problems associated with "Angrist mechanics" such as the negative intra-group correlation.[2] Further, Lavy et al. (2012) makes a convincing case that correlation due to endogenous sorting is washed out by basing conclusions only on within-individual variation in educational achievement across different learning domains (English, mathematics and science).[3] Although Lavy et al. (2012) finds no evidence of average peer effects – in contrast to the other studies in this strand of the literature – it does find that high achieving peers impact positively on low achieving girls (but not boys), and that both genders, across the ability distribution, suffer from having low achieving peers (bad apples). These effects, although statistically significant, are small in magnitude. Mendolia et al. (2018), but not Gibbons and Telhaj (2016), also finds some evidence of bad apple effects. Zhang (2016) does not examine non-linear peer effects.

3. Data and identification

The Australian testing system, NAPLAN, was introduced in 2008,

[1] Mendolia, Paloyo, and Walker (2018) takes a slightly different approach from the others in using primary school test scores of peers of peers to instrument for the secondary school test scores of peers.

[2] The negative intra-group correlation arises because of the use of the leave-out mean used in this literature, which is unnecessary in the construction of the peer effects variable from new peers only.

[3] This within-individual cross-domain approach has also been used in various studies of other education-related questions (see, for example, Altinok & Kingdon, 2012; Dee, 2005; Lavy, 2010; Schwerdt & Wuppermann, 2011).