

# Parental and School Responses to Student Performance: Evidence from School Entry Rules\*

Peter Fredriksson<sup>†</sup>      Björn Öckert<sup>‡</sup>      J. Lucas Tilley<sup>§</sup>

March 20, 2026

**Abstract:** We examine whether parental and school investments reinforce or compensate for student performance. Our analysis exploits school-starting-age rules in 34 countries, capturing achievement variation that arises because younger children typically underperform their older peers. Parents respond to lower performance by providing additional homework help and skills practice, while schools allocate weaker students to smaller classes and offer more remedial tutoring. Notably, parents provide more support to low-performing children in nearly all countries studied. Compensatory investments increase over grade levels, suggesting parents and schools respond as more information about achievement is revealed. Moreover, our evidence suggests that parental and school investments are substitutes.

**Keywords:** human capital investment, parental inputs, school inputs, student performance, school starting age

**JEL codes:** I21, I28, J24

---

\*We are thankful for helpful comments from seminar participants at the Bristol Workshop on Economic Policy Interventions and Behavior, the Norwegian School of Economics, the Swedish Institute for Social Research (SOFI), Uppsala University, Aarhus University, the Stockholm-Uppsala Education Economics Workshop, the Tinbergen Institute, SOLE, the Nordic Summer Institute in Labor Economics (Aarhus), the Research Institute of Industrial Economics (IFN), and the International Workshop on Applied Economics of Education (IWAEE). Fredriksson acknowledges funding from Forte and Handelsbanken; Tilley acknowledges funding from Handelsbanken.

<sup>†</sup>Uppsala University and Uppsala Center for Labor Studies (UCLS). Address: Department of Economics, Uppsala University, Box 513, SE-751 20 Uppsala, Sweden. E-mail: [peter.fredriksson@nek.uu.se](mailto:peter.fredriksson@nek.uu.se).

<sup>‡</sup>The Institute for Evaluation of Labour Market and Education Policy (IFAU), Uppsala University, and UCLS. Address: IFAU, Box 513, SE-751 20 Uppsala, Sweden. E-mail: [bjorn.ockert@ifau.uu.se](mailto:bjorn.ockert@ifau.uu.se).

<sup>§</sup>Uppsala University, UCLS, and Center for Educational Leadership and Excellence (CELE). Address: Department of Economics, Uppsala University, Box 513, SE-751 20 Uppsala, Sweden. E-mail: [lucas.tilley@nek.uu.se](mailto:lucas.tilley@nek.uu.se).

# 1 Introduction

Parental and school investments are crucial for children’s human capital formation.<sup>1</sup> Whether the allocation of parental and school resources reinforces or compensates for student performance thus has important consequences for the distribution of skills in both the present and future generations. Such responses also affect the efficiency of education policies aiming to reduce performance gaps and can help explain why interventions have differential effects across contexts and family background (Todd and Wolpin, 2003).

Using data from education systems around the world, this paper studies how parental and school investments adjust to student performance. We also examine how parents’ and schools’ responses vary over grade levels and whether they interact with one another. Our empirical approach leverages exogenous variation in achievement due to students’ expected relative age—i.e., their date of birth relative to the legal school entry cutoff date—and, thus, exploits the well-established performance gaps for children who start school at different ages (e.g., Fredriksson and Öckert, 2005; Bedard and Dhuey, 2006; McEwan and Shapiro, 2008; Elder and Lubotsky, 2009).

It is theoretically ambiguous how parents and schools choose to allocate resources across low- and high-performing students. The optimal allocation depends on their objectives, the properties of the human capital production function, and the budget and time constraints they face. Decreasing returns to child quality typically cause families to invest more in weaker children (e.g., Becker, 1981).<sup>2</sup> Additionally, policymakers may have redistributive goals for schools (Chambers, 1978; Thomas, 1980). The objectives of policymakers and parents must, however, be weighed against the properties of the human capital production function, which may provide incentives to allocate more resources to high-performing students (Becker and Tomes, 1976; Cunha and Heckman, 2007). Furthermore, the availability of resources may restrict decision-makers’ possibilities to invest optimally.

In the absence of clear theoretical predictions, the responsiveness of parental and school inputs to student performance remains an empirical question. Our study provides international evidence on such responses using data from the Progress in International Reading Literacy Study (PIRLS) and the Early Childhood Longitudinal Study (ECLS). We construct measures of parental inputs (e.g., help with homework and skills practice) and school resources (e.g., class size and remedial tutoring) and document responses to performance variation across 34 countries. Additionally, we shed light on the dynamics of parental and school investments over grade levels in the US. Our analysis establishes five main facts.

First, parents compensate for academic disadvantage. Low-performing students receive sig-

---

<sup>1</sup>For evidence on the importance of school resources, see, e.g., Krueger (1999), Angrist and Lavy (1999), Fredriksson et al. (2013), and Jackson et al. (2016). For evidence on the importance of family inputs, see, e.g., Björklund et al. (2006), Holmlund et al. (2011), Grönqvist et al. (2017), and Hanushek et al. (2021).

<sup>2</sup>Parents may also be averse to income inequality across children, which can reinforce compensatory behavior in terms of transfers (Becker and Tomes, 1976) or human capital investments (Behrman et al., 1982).

nificantly more help with homework than other students. Parents also spend more time practicing literacy and numeracy skills with children who perform worse in school. We find no evidence that these responses are driven by the behavior of teachers: low-performing students do not receive more homework than high-performing students.

Second, parents' responses to student performance are consistent across contexts. Parental investments are compensatory in all 34 countries we study, and in three quarters of the countries, the response is statistically significant.<sup>3</sup> Moreover, we find parental adjustment across the entire socioeconomic background distribution, although the response is larger in absolute value among high-educated parents.<sup>4</sup>

Third, schools allocate more resources to weaker students. Low-performing children are more likely to be placed in smaller classes and are more frequently given remedial tutoring.

Fourth, there is little parental and school compensation at kindergarten entry, but such behavior emerges at subsequent grade levels. This result may be surprising, given that the disadvantages faced by children who are relatively young for their grade should be apparent at school start. Indeed, some parents seem aware of these disadvantages, as reflected in our finding that they are more likely to delay school entry for children who would otherwise be among the youngest in their cohort. We conjecture that parents lack perfect information about the strength or persistence of the relationship between relative age and performance, and that there is some information revelation over time.<sup>5</sup>

Fifth, parental responses are stronger in contexts where relative age has a larger impact on student performance. This suggests that parents compensate more when schools are less successful in reducing the relative age gradient, and that parental and school inputs are substitutes.

Our study is closely related to two recent papers that investigate parental responses to variation in relative age in Spain (Berniell and Estrada, 2020) and Chile (Celhay and Gallegos, 2023).<sup>6</sup> The focus of our paper is slightly different. Although we present reduced-form estimates, we use the stipulated school entry rules as a source of plausibly exogenous variation in student performance. This allows us to study how schools and parents respond to the same change in achievement across countries and grade levels; however, it requires that relative age does not directly impact parents' and schools' investment behavior, holding child performance constant.<sup>7</sup> While recognizing that this is a strong assumption, we show that parental investments do not change when we control for absolute age.<sup>8</sup> Another concern is that parents may be aware of

---

<sup>3</sup>This differs from Celhay and Gallegos (2023). They find that parents invest more material resources in children who are relatively old for their grade in Chile, and no differences in time investments by relative age.

<sup>4</sup>This contrasts with Berniell and Estrada (2020), who find no compensation among low-educated parents.

<sup>5</sup>See Dizon-Ross (2019) for the importance (of the lack) of information for parental behavior.

<sup>6</sup>Our paper also relates to Landersø et al. (2020) and Karbownik and Özek (2023). They study family spillovers of relative age (on, e.g., sibling's school performance and mother's employment), providing indirect evidence of parents' and schools' investment behavior.

<sup>7</sup>For related evidence on responses to academic achievement, see, e.g., Yurk Quadlin (2015), Grätz and Torche (2016), Frijters et al. (2013), Nicoletti and Tonei (2020), and Fan and Porter (2020).

<sup>8</sup>We base this analysis on data from the US and Canada, which have within-country variation in the school entry rules, enabling us to hold quarter of birth constant.

relative age effects and make anticipatory investments before school entry to help offset them. Our theoretical model shows, however, that such anticipatory behavior would attenuate the magnitude of the estimates, but not reverse their sign.

Our paper also relates to the literature on parental responses to school inputs, such as school quality (Cullen et al., 2006; Pop-Eleches and Urquiola, 2013), school grants (Das et al., 2013), class size (Fredriksson et al., 2016), and teacher qualifications (Chang et al., 2022). We extend the literature by focusing on responses to student performance directly. The estimates are thus informative about the possible adjustments of parents (and schools) to interventions with a given impact on student achievement. In addition, we provide evidence on how parental and school responses interact with one another.

Our paper belongs to a small literature directly estimating how school inputs adjust to plausibly exogenous variation in student performance.<sup>9</sup> In general, parental and school responses may explain why the effects of educational interventions vary in the distribution of socioeconomic background, as well as over grade levels—i.e., whether the effects persist or fade out.<sup>10</sup> Krueger and Whitmore (2001) and Bailey et al. (2020), for instance, argue that school responses contribute to the fade-out of effects. We provide direct evidence on such responses.

The paper unfolds as follows. In the next section, we present a theoretical framework, which we use to derive a number of predictions and discuss the challenges we confront in our empirical work. Section 3 includes a description of the data, some institutional details, and our empirical strategy. Section 4 presents the main results, and Section 5 contains some robustness checks. Finally, Section 6 concludes.

## 2 Theoretical framework

We develop a simple model of human capital production and parental investment behavior to guide our empirical analysis. The model is in the spirit of Becker (1981) and Cunha and Heckman (2007), and it helps us characterize conditions when parental behavior is compensatory or reinforcing. Moreover, it facilitates discussion of a key challenge in our empirical work: specifically, the lack of data on parents' investment behavior prior to school start. The model illustrates that anticipatory behavior, where parents act before school start to offset expected

---

<sup>9</sup>Two recent studies exploit proficiency cutoffs on standardized exams to examine how schools respond to low student performance: Figlio and Özek (2024) examine a test-based remediation policy in Florida and find that students scoring below proficiency thresholds are placed in remedial classes and receive additional resources. Similarly, Berne et al. (2025) study a retention policy in Michigan and show that third graders who narrowly miss the proficiency cutoff for grade advancement receive more intensive support than those just above the threshold, even though few are actually retained. These studies complement earlier evidence on the effects of relative age on grade retention (Bedard and Dhuey, 2006; Dhuey et al., 2019) and receipt of special education (Dhuey and Lipscomb, 2010; Shapiro, 2023).

<sup>10</sup>Several papers have examined whether the effects of educational interventions persist or fade out. Lee and Loeb (1995), Currie and Thomas (2000), and Johnson and Jackson (2019) examine preschool programs; Krueger and Whitmore (2001) and Fredriksson et al. (2013) investigate class size reductions; and Jacob et al. (2010) and Chetty et al. (2014a) examine teacher quality.

disadvantages, will attenuate our estimates while leaving our qualitative conclusions intact.

## 2.1 Human capital production function

The human capital of the child ( $H$ ) is given by

$$H = g(\theta, a)h(R)$$

where  $\theta$  denotes the child's innate ability and  $a$  the relative age.  $R$  denotes overall resources ( $h_R > 0$  and  $h_{RR} \leq 0$ ), which depend on both parental and school investments. We model parental investments as a CES aggregate of investments in period 0 and period 1, where we think of  $t = 0$  as the preschool stage and  $t = 1$  as the in-school stage:

$$R_p = [\mu_0 p_0^\tau + \mu_1 p_1^\tau]^{1/\tau}$$

Here, the parameter  $\tau \leq 1$  governs the elasticity of substitution over time, and  $\mu_t$  the relative efficiency of period  $t$  investments ( $\mu_0 + \mu_1 = 1$ ). Overall resources are then given by:

$$R = [\mu_p R_p^\phi + \mu_s s^\phi]^{1/\phi}$$

where the parameter  $\phi \leq 1$  governs the elasticity of substitution across types of investments, and  $\mu_p$  ( $\mu_s$ ) denotes the relative efficiency of parental (school) investments ( $\mu_p + \mu_s = 1$ ).

The child's relative age matters for human capital because we assume children are graded on a scale (at least partially). Children who are older at the evaluation point obtain higher grades on average. Older children thus have greater progress in the education system and ultimately higher educational attainment.

## 2.2 Parental behavior

Parents have preferences defined over their own consumption ( $c^p$ ) and the long-run consumption opportunities (i.e., income:  $y = wH$ ) of their (only) child:

$$U^p = U(c^p, y) = \ln c^p + \frac{y^{1-\gamma} - 1}{1-\gamma}$$

The budget constraint is given by  $c^p = y^p - w^p(p_0 + p_1)$ . Parents thus supply a fixed amount of efficiency units of labor, they do not leave bequests, there is no discounting, and the price of investing in  $p_t$  is  $w^p$ . Parents maximize  $U^p$  subject to the budget constraint while taking school investments as given.

We begin by solving the model under perfect information. Optimal investments are deter-

mined by

$$\frac{w}{w^p} \frac{\partial H}{\partial p_1} = MRS(c^p, y) \quad (1)$$

$$\frac{\partial H}{\partial p_0} = \frac{\partial H}{\partial p_1} \quad (2)$$

where  $MRS$  denotes the marginal rate of substitution between the parents own consumption and the child's consumption opportunities. The first condition thus trades off the child's and the parent's consumption opportunities, while the second condition determines the optimal allocation of investments over time. Equation (2) dictates that the optimal allocation of parental investments over time is given by

$$\frac{p_1}{p_0} = \left( \frac{\mu_1}{\mu_0} \right)^{1/(1-\tau)} \quad (3)$$

### Compensatory behavior

A key focus of our paper is whether parents are compensatory or not. In the model, this depends on the  $MRS$ : If the marginal utility of the child's consumption opportunities declines faster than the marginal utility of their own consumption, parents exhibit compensatory behavior. Given our specification of the utility function, we thus conclude that parents are compensatory if and only if  $\gamma > 1$ .<sup>11</sup>

We also examine how parental investments are affected by school policies. That is, if schools are compensatory, how is parental behavior affected? The answer to this question depends in part on whether parental investments and school investments are complements or substitutes in the human capital production function. If parents are compensatory (i.e.,  $\gamma > 1$ ),  $\phi > 0$  unambiguously implies that parents reduce their investments when schools compensate for performance by investing more in lower-performing children.

### Parental information

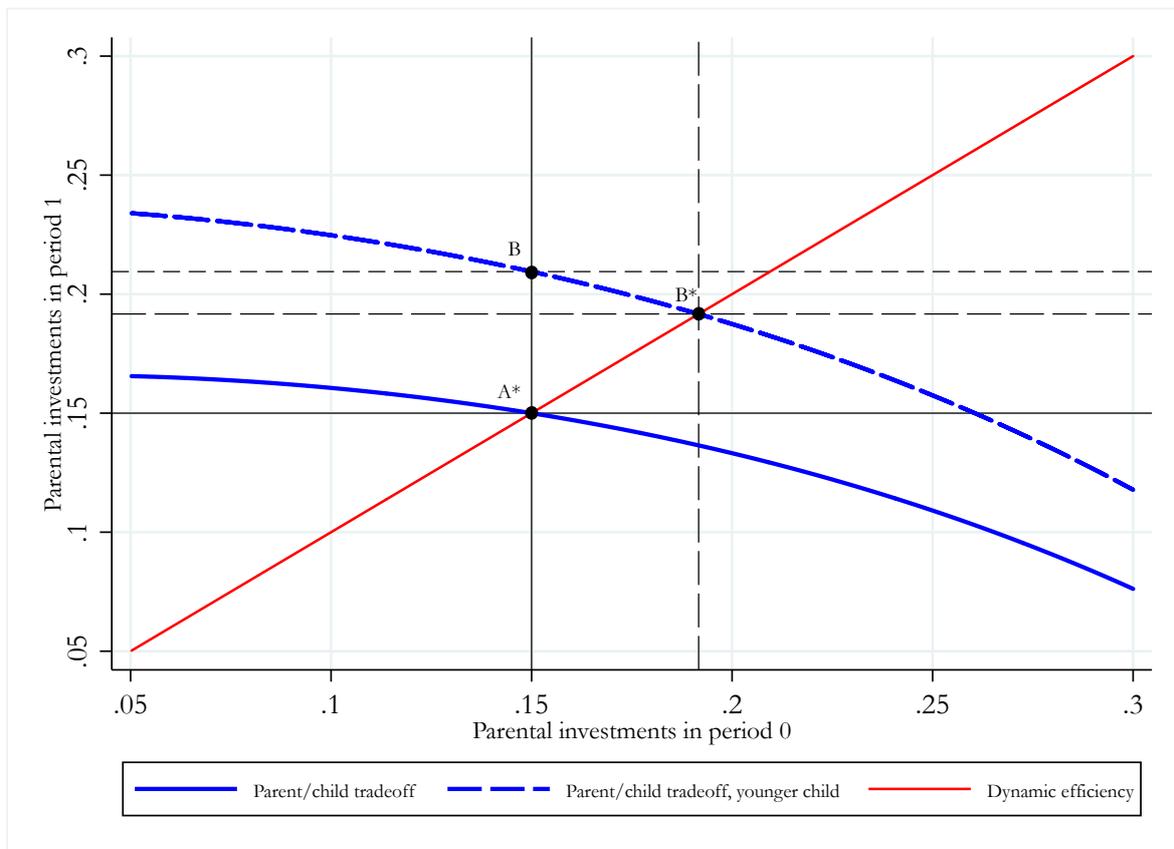
Figure 1, inter alia, illustrates the allocation of parental investments over time. The upward-sloping solid (red) line depicts the dynamic efficiency condition, i.e., equation (3). The downward-sloping solid (blue) line shows the parental-child trade-off, i.e., equation (1). The equilibrium is at point A\*. Total parental investments ( $p_0 + p_1$ ) equals 0.3. As drawn—i.e., imposing  $\tau = 0$  and  $\mu_0 = \mu_1$ )—parents allocate an equal amount to both periods,  $p_0 = p_1 = 0.15$ .

Suppose school performance drops. If parents are compensatory ( $\gamma > 1$ ), the parent-child trade-off shifts outwards, as illustrated by the dashed (blue) line. For parents who could perfectly foresee this performance drop, the new equilibrium is at point B\*. Total investments go up, and—by assumption—this increase is equally allocated across the two periods. Importantly for our empirical analysis, we underestimate the total increase in parental investments, since we

---

<sup>11</sup>The appendix contains proofs of our results.

Figure 1: Parental investments in children under different information scenarios



Notes: The solid red line shows equation (3), assuming  $\tau = 0$  and  $\mu_0 = \mu_1 = 0.5$ . The solid blue line represents equation (1), with the additional assumptions that  $\gamma = 2$ ,  $\mu_p = 0.3$ ,  $s = 0.05$ ,  $\phi = 0.1$ ,  $\rho = 1$ , and  $w = w^p = y^p = 1$ .

only observe their investments during the in-school period. Note that the estimated parental response is only attenuated: the sign still reveals whether parents' behavior is compensatory.

Now, consider the opposite extreme: suppose there is a completely idiosyncratic reduction in school performance caused by relative age. Since this reduction was unforeseen,  $p_0$  is given. Hence, uninformed parents move to the constrained equilibrium, illustrated by point B. In this special case, the increase in  $p_1$  thus captures the full response of parents to the reduction in school performance.

The contrast across the two information scenarios is also of some interest. Suppose we observe the same response to a reduction in performance during the in-school period across two groups—say high- or low educated parents for the sake of concreteness. Suppose, in addition, that evidence suggests that the high-educated invest more during the pre-school period. Since attenuation bias is stronger for the well-informed group, this suggests that high-educated parents respond more strongly to reductions in school performance (despite the fact that in-school investment behavior looks to be similar).

## 2.3 Summary of predictions and discussion

The above model is simple, and the structure is, of course, adapted to our empirical setting.<sup>12</sup> By way of summary, the model has the following predictions:

**Compensatory behavior** If (and only if)  $\gamma > 1$ , parents increase investments in their children in response to a reduction in student performance (induced by lower relative age).

**Parental compensation in more compensatory contexts** If  $\phi > 0$  and  $\gamma > 1$ , parents respond to school compensation by reducing their own investments.

**Attenuation** For parents who foresee the reduction in school performance, investment responses during the school period are attenuated measures of total investments (as long as  $\tau < 1$  and  $\mu_0 > 0$ ). Crucially, we can still infer whether parents compensate or reinforce student performance from observation of behavior during the in-school period.

**Information** Attenuation bias increases in the precision of the information signal during the pre-school period. For uninformed parents, we observe the full behavioral impact during the in-school period.

---

<sup>12</sup>The structure is broadly consistent with the empirical literature. However, some modeling choices are important. The CES structure is a convenience assumption, which is sometimes found to be too restrictive (Agostinelli and Wiswall, 2025). More importantly, we have abstracted from financial bequests. In a model with financial bequests, we would observe parents compensating for age-induced performance variation only if parental investments and age are substitutes in a production sense. While such substitution is reasonable in a model with variation in absolute age or variation across grades, we think that these considerations are less important in our empirical setting.

In the sequel, we present evidence suggesting that parents compensate for reductions in school performance (so  $\gamma > 1$ ). Moreover, such behavior is less pronounced in more compensatory school contexts, which is in line with  $\phi > 0$ .<sup>13</sup> Although, we do not observe pre-school investments, we show that all parents delay the entry of children who are expected to be young for grade, and that such behavior is particularly pronounced among the high-educated; we might thus suspect larger attenuation among the high-educated. Finally, while the evidence suggests parents have some information about the performance variation induced by relative age already prior to school entry, information appears to be incomplete. Both parental and school responses tend to grow somewhat over grades. Thus, there seems to be some information revelation after school entry.

## 2.4 Empirical strategy and implementation

We focus on the investment responses to performance variation induced by expected relative age. For given school entry age legislation, expected relative age is determined by birth timing and thus essentially randomized at birth. We identify the effects using within-school-cohort variation in relative age. This relies on stronger assumptions than the typical between-school-cohort variation used in a Regression Discontinuity framework. Nevertheless, our approach holds grade level (and time) constant. Also, our balancing exercises do not refute randomization at birth, as we illustrate in sequel.

That expected relative age is randomly assigned with respect to background characteristics does, of course, not preclude anticipatory behavior, as emphasized by the model. Nor does it imply that relative age is excludable from the investment equation conditional on student performance. While the model implies excludability by assumption—age and investments are functionally separable in human capital production—we discuss the exclusion restriction further below. In particular, we show that our results are robust to controlling for age, which (at least) does not reject excludability. Nevertheless, we present both reduced form estimates of the impact of expected relative age and Instrumental Variables (IV) estimates of the effect of student performance. The main advantage of the IV approach is that it scales the estimates such that we can compare them across countries and across grades. Such scaling is necessary because compliance with school entry rules varies across countries, and because relative age is differentially important across grades. An alternative scaling that adjusts for differential compliance is obtained using the first-stage impact on relative age. We report such estimates as well, and they do not change our overall conclusions.

A final issue relates to the measure of school performance. Ideally, we want a broad skills measure that is a reasonable proxy of human capital. Our primary proxy is reading performance, as measured by standardized literacy tests. We realize that this is an incomplete measure. Nevertheless, the results are basically invariant to using other proxies; Figure 2 shows that the expected-relative-age-gradient is very similar across a variety of human capital proxies. For the

---

<sup>13</sup>The evidence in, for example, [Fredriksson et al. \(2016\)](#) also suggests  $\phi > 0$ .

US, for example, the age gradient varies from 0.16 SD units for socio-emotional skills to 0.22 for reading literacy if we compare children who are expected to be oldest to those that are expected to be youngest in a grade.

### 3 Data, institutions, and validity of empirical strategy

We use two data sources to investigate parental and school responses to student performance. Our primary data source is the Progress in International Reading Literacy Study (PIRLS), which we supplement with the Early Childhood Longitudinal Studies – Kindergarten Cohort (ECLS-K) for the US.<sup>14</sup> The datasets provide comprehensive information on student performance and parental and school inputs, enabling us to study the allocation of inputs in different countries and, in the case of the US, over grade levels.

PIRLS is an international assessment of reading achievement for representative samples of fourth graders.<sup>15</sup> The International Association for the Evaluation of Educational Achievement (IEA) has administered the test every fifth year since 2001. In addition to the reading assessment, the data include information on parental and school inputs gathered from questionnaires sent to schools, teachers, students, and parents. Due to data constraints in several years, our analysis focuses on the PIRLS waves in 2006 and 2011. The majority of children in the dataset were born in 1995–96 or 2000–01.

ECLS-K:1999 and ECLS-K:2011 are longitudinal studies conducted by the National Center for Education Statistics (NCES) in the US. The studies follow children who entered kindergarten in the 1998/99 and 2010/11 school years, with the majority born in 1992–93 and 2004–05. Throughout the children’s elementary school years, the studies collect a wide range of data through child assessments, parent interviews, and teacher and school administrator questionnaires. These data include information on children’s reading and math proficiency, as well as their home, school, and classroom environments.

#### 3.1 Key variables

We use the PIRLS and ECLS-K data to construct four key variables: (1) expected relative age, (2) student performance, (3) parental inputs, and (4) school inputs. We provide the most relevant details regarding the construction of these variables below and offer comprehensive explanations in the data appendix. Tables A.1, A.2, and A.3 in the main appendix provide descriptive statistics for our estimation samples.

---

<sup>14</sup>Data from the PIRLS assessments and ECLS-K studies are publicly available through the [IEA Data Repository](#) and the [NCES Data Products](#).

<sup>15</sup>Several countries test students in other grade levels, typically fifth grade.

**Expected relative age** Expected relative age reflects whether students would be relatively young or old for their grade, assuming they entered school “on time” based on their birthday and the school starting age (SSA) regulations in their country. We percentile rank students by cohort and country,<sup>16</sup> with values of zero and one corresponding to those expected to be the youngest and oldest in their grade.

To measure expected relative age, we require information about the SSA rules in each country or region. Due to data restrictions in the ECLS-K public-use files, we use a data-driven approach to determine the relevant cutoff date for US students (see section B.2 in the data appendix for details). For all other countries, we obtain the SSA rules from the PIRLS Curriculum Questionnaire. We restrict our sample to countries with unambiguous cutoff dates that we can confirm in the observed distribution of birthdates. Figure A.1 in the appendix shows a map of the 34 countries that remain in our sample after this restriction.

**Student performance** We use the PIRLS and ECLS-K reading assessments as measures of student performance, standardizing the item response theory (IRT) scale scores by country, wave, and grade level.<sup>17</sup> We focus on reading proficiency because that is the most widely available performance measure in our data. Using the subset of students who also took assessments in math and science, panel (a) of Figure 2 shows that relative age has strikingly similar impacts on test scores in different subjects, suggesting that reading scores provide a suitable proxy for overall performance in our analysis.

Parents and schools do not receive the results of the PIRLS and ECLS-K assessments and thus cannot react directly to students’ scores. However, our US data indicate that these scores are highly correlated with other performance measures that should be more salient to parents and schools, such as teacher evaluations. For example, panel (b) of Figure 2 shows that the relative age gradient for teacher evaluations of children’s literacy skills is 0.176 SD compared to 0.217 SD for the ECLS-K standardized reading assessment. This panel also shows that the age gradient for students’ socio-emotional skills—as assessed by the teacher—is of a similar magnitude (0.155).

**Parental inputs** We obtain information on parental inputs from the home survey (PIRLS) and parent interviews (ECLS-K), which are typically completed by the child’s mother or female guardian. Due to data availability, our main measure of parental investment is the frequency with which parents help their children with schoolwork or homework<sup>18</sup>—arguably one of the most crucial school-related activities in which parents can directly participate. We standardize

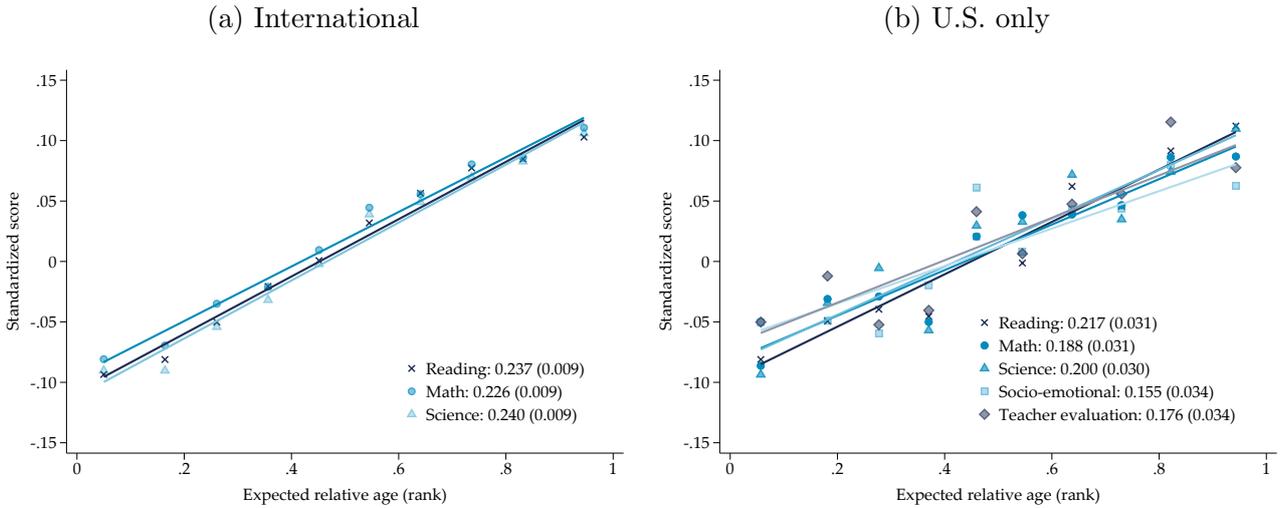
---

<sup>16</sup>In countries where the SSA rule varies across subnational jurisdictions, we compute the rank at the relevant regional level rather than the country level.

<sup>17</sup>The PIRLS data includes five plausible values for individual reading scores. We use the first value for our main analysis, and as a robustness check, verify that the results are similar using the other four values.

<sup>18</sup>The exact phrasing of the question varies somewhat by data source and survey wave. See section B.3.4 in the data appendix for details.

Figure 2: Relative age gradients by skills domain



*Notes:* Each line depicts the relationship between children’s expected relative age and their standardized score in a specific skill domain. Panel (a) uses international data from PIRLS 2011, TIMSS 2011, and the fifth-grade waves of ECLS-K:1999 and ECLS-K:2011, covering 115,902 students from 22 countries. We plot separate lines for standardized assessments in reading, math, and science, with the estimated gradients and standard errors reported in the legend. Panel (b) presents the same analysis for US data from the ECLS-K studies ( $N = 10,550$  students), along with additional lines for teacher evaluations of children’s literacy skills and socio-emotional skills. See the data appendix for details about data measurement.

parents’ responses by country, wave, and grade level. We also create standardized measures of other school-related parental time investments, such as the frequency with which parents practice basic skills with their children.

**School inputs** We obtain information on school inputs from teacher questionnaires. Class size is our main measure of school inputs. For the US, we have richer data on school inputs because teachers complete child-level questionnaires in addition to providing information at the class level. When analyzing US data, we also use an indicator of whether the child receives remedial tutoring at school.

### 3.2 Empirical strategy

We study how parents and schools adjust their investments in children in response to variation in their academic performance. The naive OLS estimate of parental (or school) investments on children’s performance is biased due to omitted variables, reverse causality, or measurement errors. To address the bias, we implement a two-stage least squares (2SLS) approach, exploiting the well-known fact that children who are relatively old for their school cohort tend to outperform their younger peers (e.g., Fredriksson and Öckert, 2005; Bedard and Dhuey, 2006; McEwan and Shapiro, 2008; Elder and Lubotsky, 2009).

Specifically, our approach leverages performance variation stemming from children’s date of birth relative to the school starting age rule in their country or region. We estimate first-stage

regressions of the following form:

$$H_{ij} = \alpha_j + \eta a_{ij} + \epsilon_{ij}, \quad (4)$$

where  $i$  indexes individuals and  $j$  schools. We acknowledge that the impact of age on achievement ( $\eta$ ) is a reduced-form (RF) parameter that partly reflects compensatory behavior among parents and schools. Moreover,  $\eta$  varies over countries because school starting age rules are differentially predictive, e.g., due to cross-country variation in the extent of “red-shirting” (delayed school entry).

The reduced-form relationship between, say, parental investments,  $p_{ij}$ , and  $a_{ij}$  is given by

$$p_{ij} = \delta_j^p + \pi^p a_{ij} + \omega_{ij}^p \quad (5)$$

The instrumental variables (IV) estimate of the response of parental investments to performance variation is thus obtained from equations (4) and (5) as:  $\beta^p = \pi^p / \eta$ . Analogously, the IV estimate of the response of school resources,  $s_{ij}$ , to performance is given by  $\beta^s = \pi^s / \eta$ , where  $\pi^s$  is the coefficient in the reduced-form relationship between school resources and the expected relative age rank, and  $\eta$  comes from equation (4). We present estimates for single countries as well as estimates that pool the data for all countries to estimate an average response to changes in expected relative age. The IV approach achieves comparability across countries and grade levels by estimating how parents and schools respond to the same change in achievement.<sup>19</sup>

### 3.3 Validity of the empirical strategy

Strictly speaking, the instrumental variable approach requires: relevance, randomization, excludability, and monotonicity. Regarding relevance, panel (a) of Figure 3 provides a graphical depiction of the first-stage regression, revealing a strong relationship between expected relative age and standardized reading score (t-ratio = 33.64). On average, the oldest children in a grade score almost a quarter of a standard deviation higher than their youngest peers.

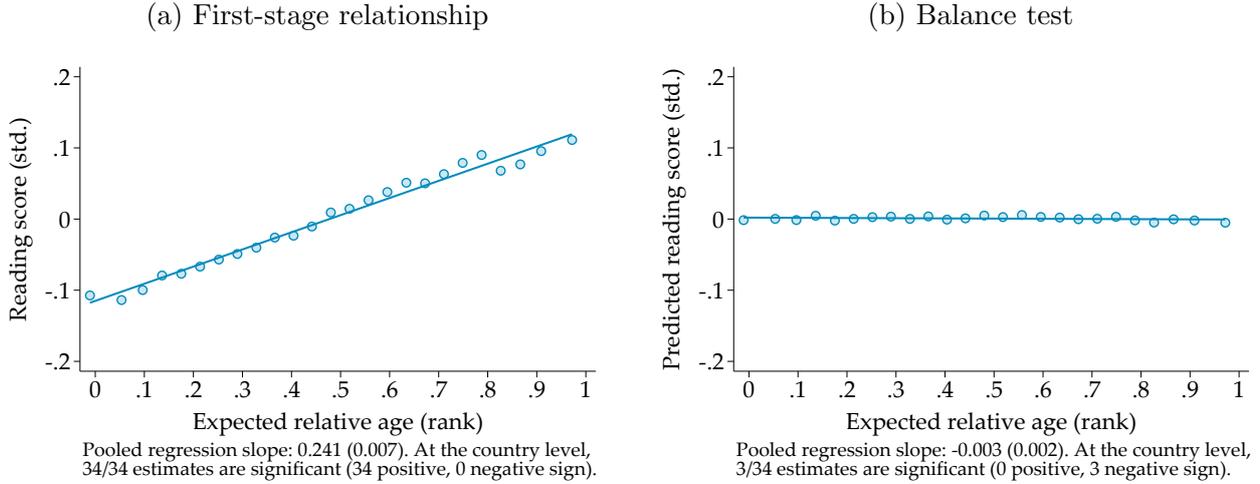
Regarding randomization, expected relative age is likely as good as randomly assigned at birth. Panel (b) of Figure 3 shows a balance test for predicted reading score, indicating that children’s expected relative age is unrelated to pre-determined background variables (gender, foreign background, and parental education). This result suggests that children born in different parts of the school year are comparable, thus validating that birth timing is as good as randomly assigned relative to the school-entry cutoffs.<sup>20</sup> A challenge is that parents may preemptively invest in their children prior to school start. Our theory suggests, however, that such behavior

---

<sup>19</sup>Some of the variation across countries may, of course, come from the complier populations being different across countries.

<sup>20</sup>When running country-specific regressions, three nations fail the balance test: Croatia, Hong Kong, and Singapore. In the appendix, we show that our main estimates are insensitive to dropping these three nations (see Table A.5).

Figure 3: Validity of the empirical strategy



*Notes:* Panel (a) shows the first-stage relationship from equation (4), and Panel (b) shows the relationship between the predicted reading score and the expected relative age rank. The prediction uses gender, foreign background, and parents' highest level of education as predictor variables. In both panels, we use pooled international data from PIRLS 2006, PIRLS 2011, and the fifth grade waves of ECLS-K:1999 and ECLS-K:2011 ( $N = 287,675$ ). We residualize on school fixed effects before binning and plotting.

yields attenuated estimates of the relationship between parental investments and achievement using data measured after school start.

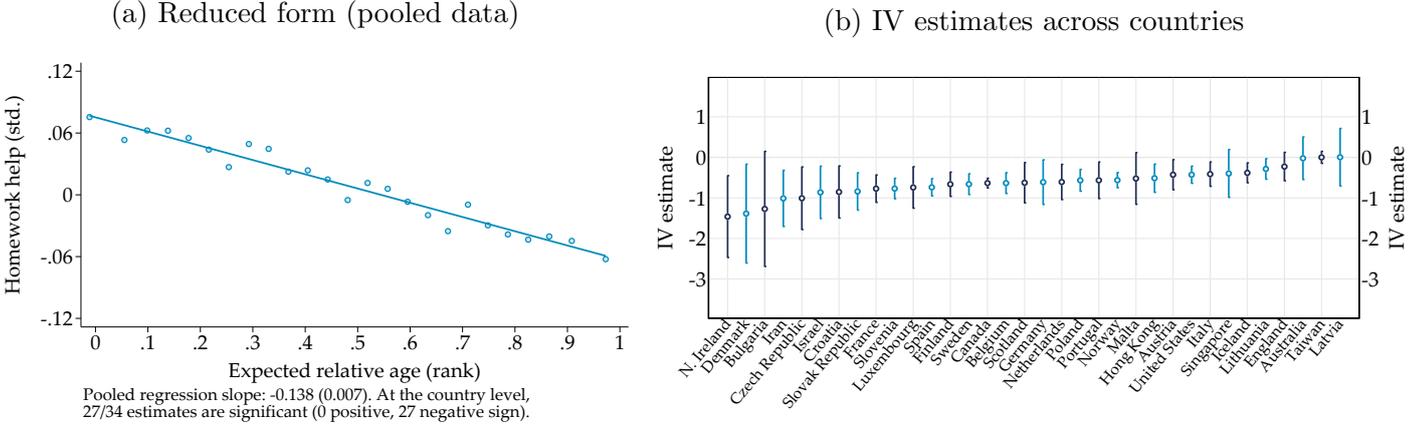
The exclusion restriction—i.e., that parents and schools respond to performance rather than directly to children's expected relative age—may also be a concern. In the sequel, we shed light on the validity of this assumption by showing that our estimates of parental responses are largely unaffected by controlling for absolute age.

## 4 Results

### 4.1 Parental responses: Help with homework

We begin by documenting that parents make compensatory investments in their children's learning. Figure 4a depicts the reduced-form relationship between a child's expected age rank and the amount of homework help that parents provide. On average, the youngest children in a grade receive 0.14 standard deviations more homework help than the oldest children. Scaling this number by the corresponding gap in reading performance yields an implied IV estimate of  $\hat{\beta}^p = -0.14/0.24 = -0.58$ ; see Table 1 below. Thus, when children's reading performance decreases by one standard deviation, parents increase their homework help by 0.58 standard deviations. This amounts to about 1.19 additional days per week on a base of 2.77 days. The observed response does not arise because teachers support weaker students by assigning them extra homework; Table A.6 shows that the youngest and oldest children in a grade receive similar workloads.

Figure 4: Parental help with homework



*Notes:* Panel (a) illustrates the reduced-form relationship between parental homework help and expected relative age using pooled data from PIRLS 2006, PIRLS 2011, ECLS-K:1999 (grade 5), and ECLS-K:2011 (grade 5). We have residualized on school fixed effects before binning and plotting. Panel (b) plots country-specific estimates from instrumental variables specifications in which we use expected relative age as an instrument for standardized reading performance. All regressions include school fixed effects.

Table 1: Parental responses to variation in children’s reading proficiency

	Panel (a): Full sample			Panel (b): US and Canada		
	(1)	(2)	(3)	(4)	(5)	(6)
Reading score (std.)	-0.238 (0.004)	-0.579 (0.029)	-0.547 (0.028)	-0.567 (0.055)	-0.529 (0.052)	-0.749 (0.195)
Estimation type	OLS	IV	IV	IV	IV	IV
Controls:						
School fixed effects	✓	✓	–	✓	–	–
Country fixed effects	–	–	✓	–	✓	✓
Database fixed effects	–	–	✓	–	✓	✓
Expected absolute age (linear)	–	–	–	–	–	✓

*Notes:* We estimate all regressions using pooled data from four databases: PIRLS 2006, PIRLS 2011, ECLS-K:1999 (grade 5), and ECLS-K:2011 (grade 5). The dependent variable is the frequency with which parents help their children with homework (standardized by database, country, and year). The number of observations is 251,596 for the full sample and 48,918 for the US and Canada. Standard errors are shown in parentheses and clustered by expected date of birth.

Panel (a) of Table 1 compares the OLS estimate of the relationship between parental responses and the reading proficiency of the child (column 1) with our main IV estimate (column 2). The OLS estimate is smaller in absolute value than the corresponding IV estimate. Parental compensation is thus stronger than revealed by the descriptive relationship, possibly because

highly skilled parents (in the unobserved sense) tend to invest more in their children and have children who do better in school. Furthermore, column (3) shows that the baseline IV estimate is invariant to the use of country fixed effects rather than school fixed effects. This finding is reassuring because one concern could be that we are removing an important dimension of parental investment behavior by including school fixed effects.

Not only is parental compensation strong on average, compensation is also pervasive across countries as well as across the distribution of family background. Figure 4b presents country-specific IV estimates for parents' help with homework.<sup>21</sup> There is significant compensatory behavior in three quarters of the countries in our data (p-value < 0.05), and there is no country where parents provide less homework help when their children's reading performance declines.<sup>22</sup> Furthermore, Table 2 shows that parents respond similarly to changes in performance for students with different characteristics. Columns (1)–(4) reveal no significant differences by foreign background or gender of the child, although the estimates suggest slightly stronger responses among natives.

Table 2: Heterogeneity in parental help with homework (IV estimates)

	Gender		Foreign background		Parental education	
	Girl (1)	Boy (2)	Native (3)	Foreign (4)	High (5)	Low (6)
Reading score (std.)	−0.618 (0.041)	−0.560 (0.048)	−0.602 (0.029)	−0.453 (0.089)	−0.704 (0.053)	−0.531 (0.034)
First-stage estimate	0.245	0.239	0.246	0.232	0.247	0.248
Mean homework help (std.)	−0.023	0.040	0.011	−0.032	−0.098	0.053
School fixed effects	✓	✓	✓	✓	✓	✓

*Notes:* All regressions are estimated using pooled data from PIRLS 2006, PIRLS 2011, ECLS-K:1999 (grade 5), and ECLS-K:2011 (grade 5). Parents are considered highly educated if they have a university degree. Students are considered to have a foreign background if they are exposed to a foreign language at home. The regression samples consist of 125,441 girls and 125,695 boys; 202,238 and 35,840 children with native and foreign backgrounds; and 75,238 and 164,055 children with high- and low-educated parents. Standard errors are shown in parentheses and clustered by expected date of birth.

We find some evidence of heterogeneous compensatory behavior by students' socioeconomic background. Columns (5)–(6) show that highly educated parents respond more strongly to their children's academic performance than less-educated parents, although the difference is notably smaller compared to findings in other studies (e.g., [Berniell and Estrada, 2020](#); [Fredriksson et al., 2016](#)). With that said, high-educated parents may be particularly prone to invest in their

<sup>21</sup>For the sake of comparability, we only present IV estimates. There is substantial variability in first-stage estimates across countries, e.g., due to different grade retention policies or differential enforcement of school entry rules. See Table A.7 for additional country-specific estimates.

<sup>22</sup>We have examined whether the differential responses across countries are due to differences in tracking policies across countries. We do not find any such differences.

children prior to school start, suggesting more attenuation among the high-educated than among the low-educated. Consistent with this, Figure 7 (below) shows that high-educated parents are more likely to delay the entry of children who are expected to be young-for-grade.

The IV estimates rely on the assumption that parents react to changes in their children’s performance rather than directly to their (relative or absolute) age. If this assumption holds, our IV approach should result in similar estimates when controlling for children’s (expected) age. Estimating such a model poses an empirical challenge because it requires independent variation in relative age for a given absolute age. Most countries in our data have a universal school entry rule, however, and children born on the same date thus have identical relative ages. Furthermore, in the few countries where SSA rules vary across jurisdictions, most cutoff dates are only a few months apart. The absolute and relative age are therefore highly correlated, making it hard to disentangle the two effects.

With the aforementioned caveat in mind, panel (b) of Table 1 examines the plausibility of our identifying assumption by exploiting within-country variation in school entry rules in the US and Canada.<sup>23</sup> Intuitively, we compare parental responses for children who are born at the same point in time but have different relative age ranks as a consequence of their local SSA rule.<sup>24</sup> To accomplish this, we replace the school fixed effects in our main model with country fixed effects, such that our identifying variation stems from differences in SSA rules across regions. Additionally, we include a linear control for expected absolute age.<sup>25</sup>

Column (4) of Table 1 shows that the IV estimate for the US and Canada is basically identical to the estimate for the full sample (c.f. column 2). The IV estimate does not change substantively when using country rather than school fixed effects (columns 4 vs. 5) or when adding controls for absolute age (column 5 vs. 6).<sup>26</sup> We thus view the evidence in Table 1 as not refuting the exclusion restriction.

An alternative IV approach uses the first-stage relationship between relative age and expected relative age to rescale the reduced form estimates. By doing so, we adjust for differential compliance rates with school entry rules and differential retention rates across countries. Using this alternative IV approach, we conclude that parents compensate for age and thus invest more in younger children: the IV estimate is  $-0.20$  (SE: 0.01).<sup>27</sup> Our preferred interpretation is that such compensatory behavior arises because younger children are disadvantaged in school.

---

<sup>23</sup>In principle, there is variation in school entry age rules in Australia as well. However, 50 percent of parents do not complete the questionnaires, rendering these data less useful.

<sup>24</sup>For Canada, the rules vary at the provincial level. For the US, they vary by state or local education authority. See the data appendix for details on how we determine the relevant SSA rule. Since the school year starts around the same time across jurisdictions and we compare students in the same grade level, children should have the same amount of schooling at the test date, even though the SSA regulations differ.

<sup>25</sup>We use expected absolute age rather than actual absolute age so that the results are not driven by children who repeated a grade, skipped a grade, or did not comply with the SSA rule at the time of primary school entry.

<sup>26</sup>If we control for expected age linearly, rather than using quarter-by-year-of-birth fixed effects, we find an IV estimate of  $-0.748$  (SE: 0.195), using a specification which otherwise corresponds to column (6) of Table 1.

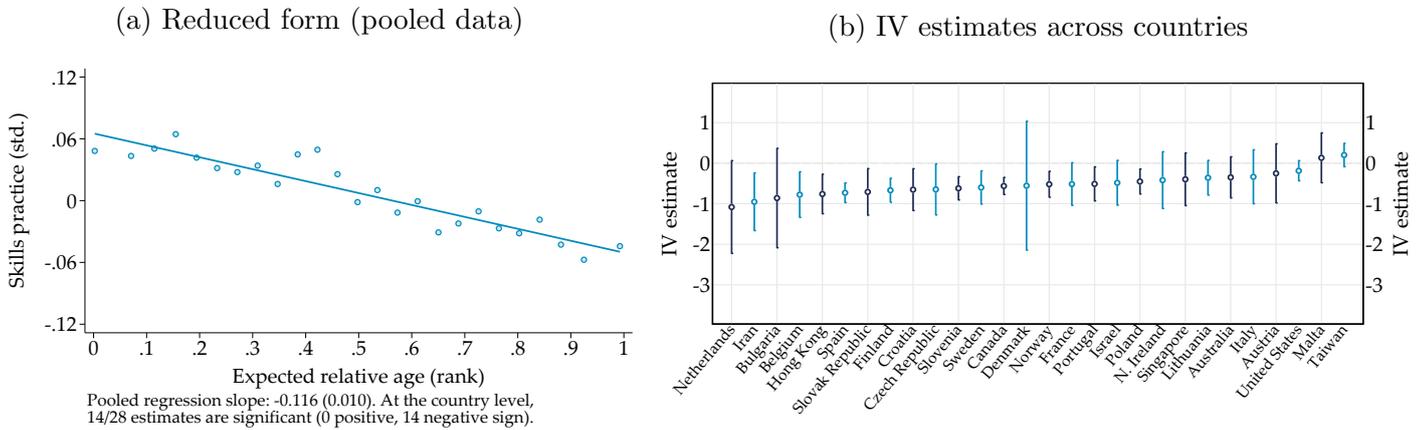
<sup>27</sup>Figure A.2 shows separate IV estimates by country.

## 4.2 Other parental responses

A possible concern with parental help with homework is that it might reflect struggling students requiring more help rather than the investment behavior on the part of parents. In this section, we thus look at other parental responses that may be more closely tied to choices made by parents.

Figure 5 examines how parental skills practice changes in response to performance variation generated by expected relative age. The information on skills practice is available for fewer countries, and only for PIRLS 2011 and ECLS-K 2011. Nevertheless, the estimates are very similar to the ones reported for parental help with homework. The reduced form estimate, illustrated in panel (a), is  $-0.116$  and the corresponding IV estimate is  $-0.503$  (SE: 0.041).

Figure 5: Parental skills practice



*Notes:* Panel (a) illustrates the reduced-form relationship between parental skills practice and expected relative age using pooled data from PIRLS 2011 and ECLS-K:2011 (grade 5). We have residualized on school fixed effects before binning and plotting. Panel (b) plots country-specific estimates from instrumental variables specifications where expected relative age is used as an instrument for standardized reading performance. All regressions include school fixed effects.

Moreover, panels (c) through (e) of Table A.8 show that our findings hold when studying other dimensions of parental time investment. We find significant compensatory responses for homework monitoring (panel c: IV-estimate,  $-0.373$ ), discussion about schoolwork (panel d: IV-estimate,  $-0.126$ ), and discussion about reading (panel e: IV-estimate,  $-0.116$ ). The evidence thus suggests that parents make compensatory investments along a variety of dimensions.

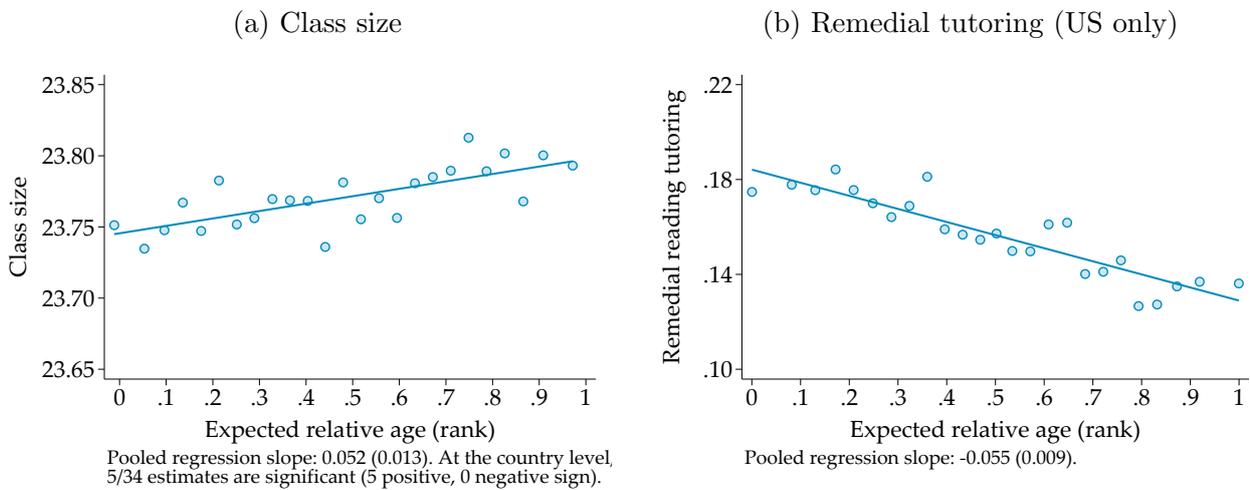
## 4.3 School responses

This section turns to the question of how schools respond to variation in student performance. Conceptually, we think of schools as operating under a fixed budget constraint, and thus, we examine whether there is compensatory or reinforcing resource allocation across students within

schools.<sup>28</sup> To align with this conceptual benchmark, we include school fixed effects in our empirical estimations. Note, however, that including school fixed effects does not affect the point estimates; their inclusion only makes the estimates more precise.

Figure 6 shows that schools make compensatory investments in children. Panel (a) reveals that class size increases with children’s expected relative age rank. The implied IV estimate equals  $\hat{\beta}^s = 0.05/0.24 = 0.22$ , indicating that schools reduce class size by 0.22 students when achievement declines by one SD. This corresponds to about 12 percent of the within-school standard deviation in class size.<sup>29</sup>

Figure 6: School responses – reduced forms



*Notes:* Panel (a) depicts the reduced-form relationship between class size and expected relative age rank using pooled international data from PIRLS 2006, PIRLS 2011, ECLS-K:1999 (grade 5), and ECLS-K:2011 (grade 5). Panel (b) depicts the reduced-form relationship between remedial tutoring and expected relative age rank using US-based data from the fifth-grade wave of ECLS-K:1999 and ECLS-K:2011. In both panels, we have residualized on school fixed effects before binning and plotting.

Panel (b) of Figure 6 illustrates the reduced-form estimate for remedial tutoring. At this stage of the analysis, we restrict attention to the US because information on this outcome does not exist in the international data. The figure shows that children who are young for their grade are assigned to more in-school tutoring. The implied IV estimate indicates that a reduction in student performance by one SD increases the probability that children receive remedial tutoring by 24 percentage points (from a mean of 16 percentage points).

Table 3 summarizes the main results for school responses. Columns (2) and (4) present IV estimates of the relationship between school responses and the reading proficiency of the child for the full sample and the US, respectively. For comparison, columns (1) and (3) show

<sup>28</sup>In the appendix, we present a model of school behavior.

<sup>29</sup>In many countries in our data, there is little to no variation in class size due to the sampling design in PIRLS. This limits our ability to compare estimates across countries. Note that the estimate for the US, where there is significant within school variation, is much larger. Class size increases one-for-one with pupil achievement in the US; see Table 3 and Figure 8.

the corresponding OLS estimates. The OLS estimates are uniformly smaller in absolute value compared to the IV estimates. School compensation is thus stronger than revealed by the descriptive relationship possibly because highly skilled parents (in the unobserved sense) can circumvent, e.g., class size allocations, at least to some extent.

Table 3: School responses to variation in child reading proficiency

	Panel (a): Full sample		Panel (b): US only			
	Class size		Class size		Remedial tutoring	
	(1)	(2)	(3)	(4)	(5)	(6)
Reading score (std.)	0.134 (0.010)	0.218 (0.053)	0.871 (0.040)	1.087 (0.399)	-0.140 (0.004)	-0.237 (0.036)
Estimation type	OLS	IV	OLS	IV	OLS	IV
School fixed effects	✓	✓	✓	✓	✓	✓

*Notes:* The table reports OLS estimates from regressions of each school input on children’s reading performance, as well as IV estimates from 2SLS regressions using children’s expected relative age rank as an instrument for reading performance. Standard errors are shown in parentheses and clustered by expected date of birth. In panel (a), the international data come from PIRLS 2006, PIRLS 2011, ECLS-K:1999 (grade 5), and ECLS-K:2011 (grade 5). In panel (b), the US data come from the fifth-grade waves of ECLS-K:1999 and ECLS-K:2011. The number of observations is 271,690 for class size (pooled), 20,631 for class size (US), and 20,683 for remedial tutoring (US). Average class size in the pooled sample (US sample) is 23.772 (23.041) with a standard deviation of 6.098 (5.845). The share of children who receive remedial tutoring is 15.655%.

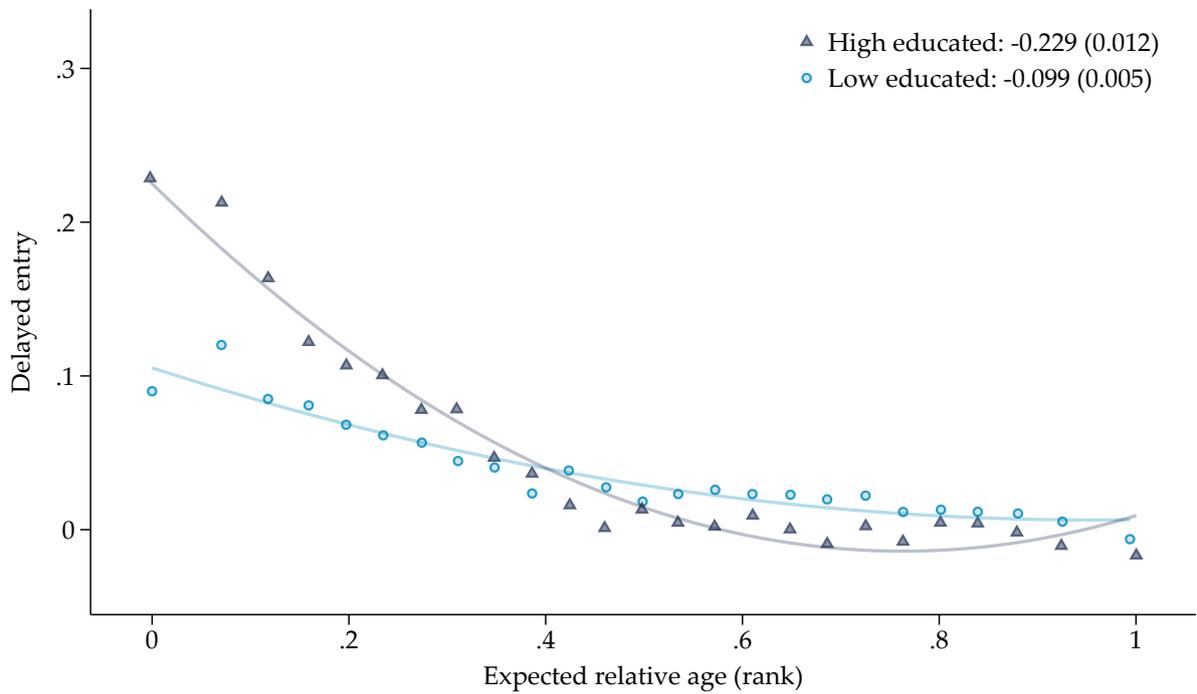
## 4.4 Responses over grade levels

An interesting question is how parental and school responses evolve over grade levels. Moreover, parents might act preemptively to counter the predicted disadvantage of the youngest children in the class. In this section, we shed light on these questions by using longitudinal data from the ECLS-K studies.

Figure 7 sheds light on preemptive behavior by examining whether delayed kindergarten entry is affected by expected relative age. It shows that children who are expected to be the youngest in their school cohort are considerably more likely to have delayed kindergarten entry compared to their older peers. Such redshirting behavior is also substantially more common among high-educated parents, suggesting that they are better informed about the importance of relative age for school outcomes.

Figure 8 turns to the question of how parental and school responses evolve over grades. The top panel pertains to help with homework; in addition, we show how parents respond in terms of practicing basic skills (reading, writing, and numeracy) with their children—an outcome we observe from the start of kindergarten. The second panel pertains to school responses, focusing on class size and remedial tutoring (note that we have inverted the scale for class size to improve readability).

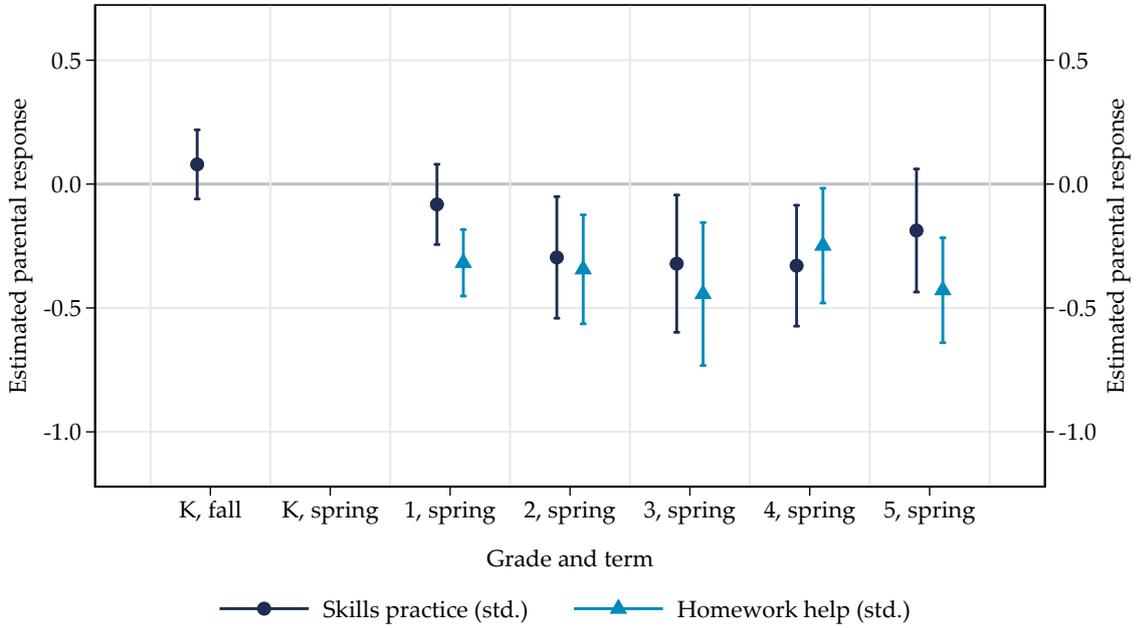
Figure 7: Redshirting by parental education



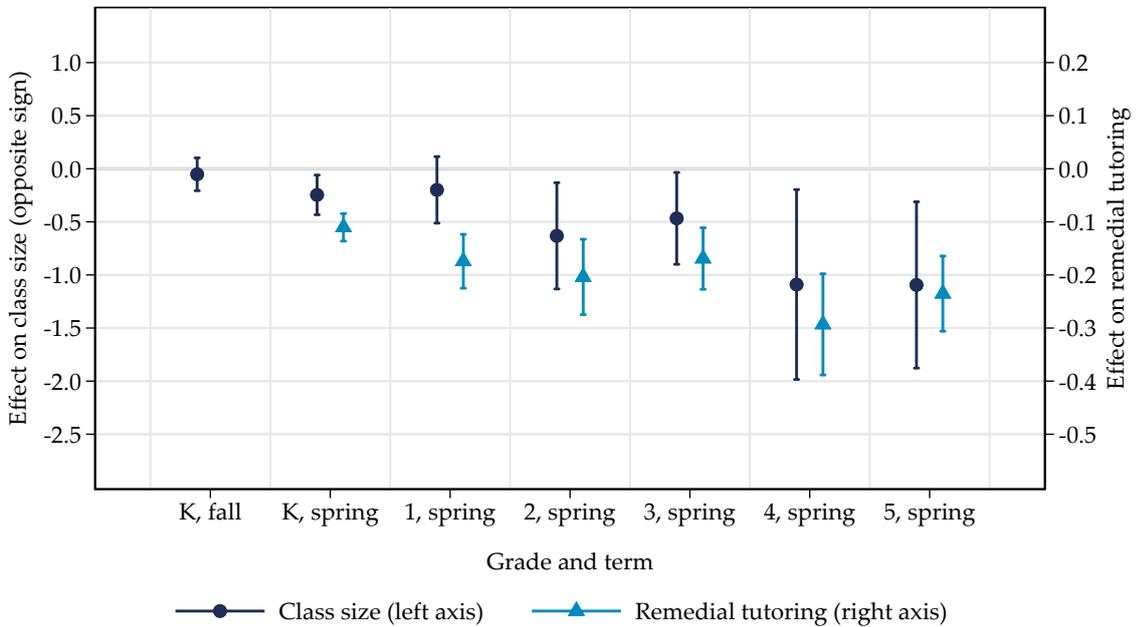
*Notes:* We measure delayed kindergarten entry (i.e., redshirting) by comparing a child's expected date of birth (assuming they adhered to the school-starting-age rule in their region) to their actual date of birth. The coefficient reported in the legend reports the linear relationship between delayed entry and relative age. Data come from ECLS-K:1999 and ECLS-K:2011.

Figure 8: Parental and school responses over grade levels in the US

(a) Parental responses



(b) School responses



Notes: This figure shows the IV estimates over grade levels using pooled data from ECLS-K:1999 and ECLS-K:2011. All regressions include fixed effects for children's base-year school. The bars indicate 95% confidence intervals when standard errors are clustered by expected date of birth. To enhance readability, we have reversed the sign of the class size estimate.

Panel (a) shows that parents do not respond at kindergarten entry: the estimate is not statistically significant and has an unexpected sign. After kindergarten, however, there is evidence of compensatory behavior. For instance, in the spring of second grade, parents respond to a one SD reduction in reading achievement by increasing basic skills practice, as well as homework help, by 0.3 of a standard deviation. Whereas parental responses are fairly stable throughout the remainder of primary school, panel (b) of Figure 8 demonstrates that the compensatory responses of schools grow over grades—from zero at kindergarten entry for class size to being quite substantive by the end of primary school. In fifth grade, a one SD reduction in performance leads to a 1.1 student reduction in class size and a 0.24 increase in the probability of receiving remedial tutoring.

What explains the variation over grades? One contributory factor may be that parents (and schools) do not have precise information about reading achievement in kindergarten. Indeed, previous research has shown that parents are partially uninformed about the progress of their children and that they change their investment behavior when new information arrives; see [Dizon-Ross \(2019\)](#). To shed light on this explanation, we analyze whether parents and teachers adjust their expectations about children’s future educational attainment as children progress through primary school—i.e., as more information on their performance is revealed.

For each grade level, the darker estimates in Figure 9 show the reduced-form relationship between children’s relative age rank and parents’ expectations regarding college completion. The estimated relationship grows stronger over time. There is an insignificant relative age gradient in kindergarten, but by fifth grade, parents of the oldest children are four percentage points more likely than the parents of the youngest children to believe that their child will obtain a college degree.<sup>30</sup> The observed pattern is consistent with parents updating their information on children’s performance over time in primary school.

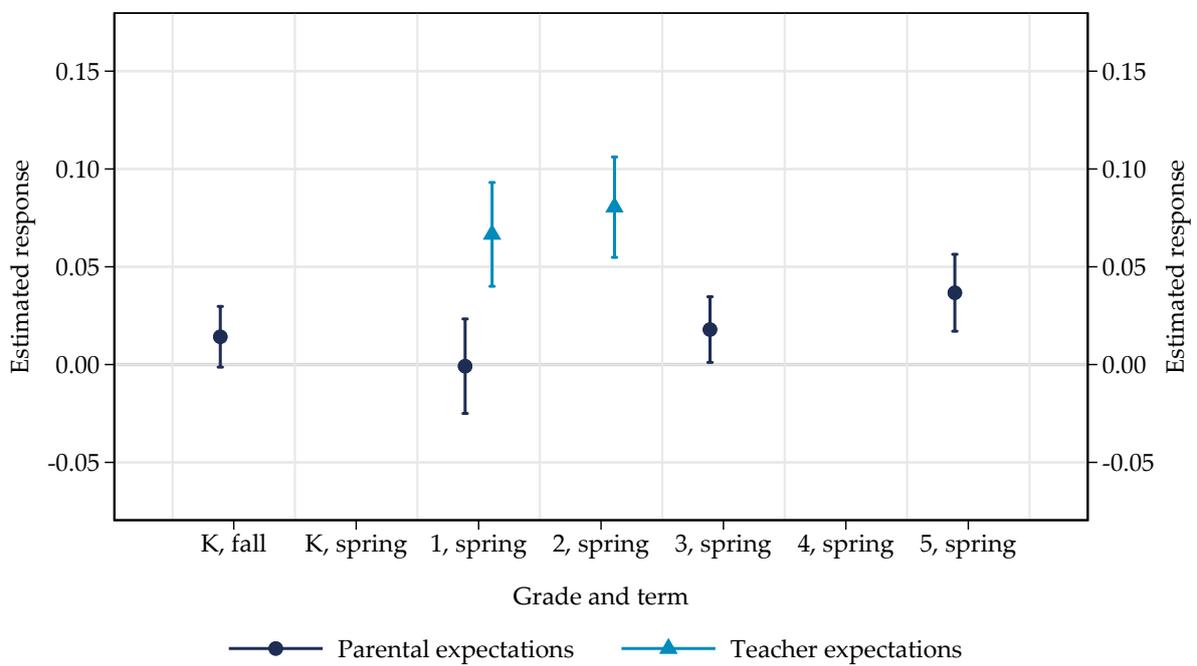
For completeness, Figure 9 also shows the reduced-form relationship between children’s relative age and teachers’ educational expectations. Data is available only for first and second graders, making it difficult to assess how expectations evolve as students age. Interpreted literally, however, it seems that teachers are better informed than parents about the impact of relative age on performance.

To summarize, our evidence suggests that parents have information about predicted achievement prior to school entry. Nevertheless, the information was not complete. The increasing responses over grade levels imply that more precise information about student achievement gets revealed over grade levels.

---

<sup>30</sup>In the appendix, we show that expected relative age has a positive impact on parents’ expectations in fourth grade across the countries in PIRLS; see Figure A.4.

Figure 9: Educational expectations over grade levels in the US



*Notes:* This figure shows the reduced-form estimates by grade level for parental and teacher expectations regarding children's likelihood of completing a college education. All regressions include fixed effects for children's base-year school. The bars correspond to 95% confidence intervals when standard errors are clustered by expected date of birth. Data for parental expectations come from ECLS-K:1999 and ECLS-K:2011. Data for teacher expectations come from ECLS-K:2011.

## 4.5 Parental responses in different contexts

Intuitively, we might expect that parents are more responsive to student achievement in contexts where schools are less compensatory. This intuition is based on the assumption that public and private investments in children’s learning are substitutes, which is what most previous research suggests; see, for example, [Fredriksson et al. \(2016\)](#), [Pop-Eleches and Urquiola \(2013\)](#), and [Cullen et al. \(2006\)](#).

The challenge in testing this prediction is that we lack a comprehensive measurement of how compensatory schools are. To address this issue, we implement an approach akin to that of [Chetty et al. \(2014b\)](#) and [Rothstein \(2019\)](#). We use regional differences in the relationship between children’s reading score and relative age rank to measure how compensatory schools are in a particular region. This approach relies on the premise that schools’ compensatory investments reduce the performance gap between younger and older students, and thus, the reading gradient should be smaller in regions where schools are more compensatory.

We implement our method in three steps. First, we divide each country in our data into groups of urban and non-urban schools, hereafter referred to as regions.<sup>31</sup> Next, we estimate region-specific regressions for school compensation (measured by the reading gradient in equation 4) and parental compensation (measured by the homework help gradient in equation 5). All regressions include school fixed effects to account for parental sorting across schools and urban areas. Finally, we correlate the resulting estimates, weighting by the inverse standard error of the reading gradient.

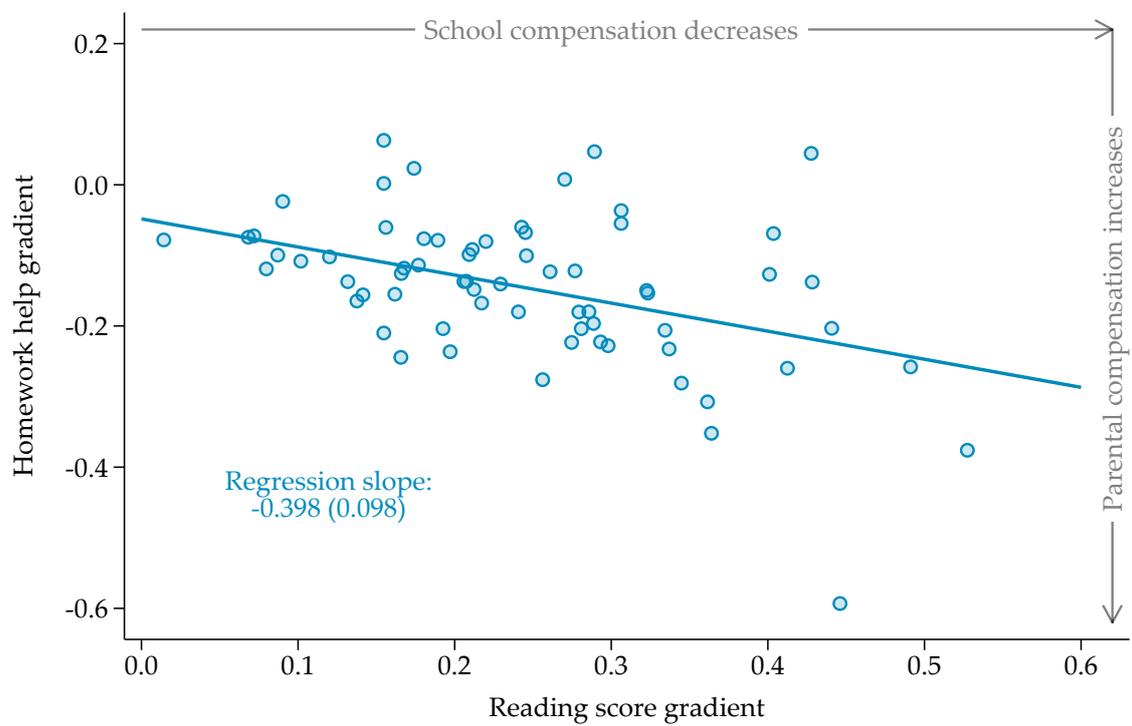
Figure 10 depicts the correlation between parental and school compensation across regions. Each point plots the reduced-form estimates for one regional subdivision.<sup>32</sup> The degree of school compensation decreases as one moves outward along the horizontal axis, whereas the degree of parental compensation increases as one moves down the vertical axis. The observed negative relationship, with an estimated slope of  $-0.40$ , thus suggests that parents compensate more in regions where schools compensate less. This finding aligns with the prediction that there is less parental compensation in more compensatory school contexts, and provides suggestive evidence that parental and school investments are substitutes rather than complements.

---

<sup>31</sup>We create regions using an indicator of school urbanicity because the data from most countries does not contain identifiers for administrative divisions, such as municipalities or provinces.

<sup>32</sup>We present robustness checks for this analysis in Figure A.5. Panel (a) shows that we obtain similar estimates if we drop the three nations that fail the balance test. Another concern is that the observed relationship between the reduced-form estimates may partly reflect differential compliance with school starting age rules across regions. Therefore, we have also estimated IV specifications in which we instrument children’s actual age rank with their expected age rank; see panel (b). We find an estimate of  $-0.32$ , which confirms the negative relationship between parental and school compensation.

Figure 10: Parental responses in different contexts



*Notes:* Each dot shows the reduced-form estimates for one region, defined as a combination of country x school urbanicity. There are 67 observations – two per country, except in the case of Singapore, where there are no rural schools. When regressing the regional homework help gradient on the regional reading gradient, each observation is weighted by the inverse standard error of the reading score gradient in that region.

## 5 Robustness checks

Figure A.6 provides several sensitivity checks of our main results for parental homework help. For instance, it shows that it does not matter which of the five plausible values we use for the PIRLS reading score, or whether we give countries equal weight in the estimations: these changes have a negligible effect on our estimates and do not affect our main conclusions.

Figure A.8 provides evidence that our results on parental responses are not driven by our focus on help with reading homework rather than other important subjects, such as mathematics. The analysis in the figure is based on the subsample of data in which parents provide answers about the extent to which they help with homework or skills practice in a specific subject (e.g., help with reading vs. math homework or practicing reading vs. math skills). Using data from the US, the left side of the figure shows that the IV estimates for help with math homework are similar—albeit slightly larger in absolute magnitude—than those for help with reading homework. The right side of the figure uses international data and also reveals somewhat stronger, though not statistically different, compensatory responses for skills practice in math compared to skills practice in reading.

Finally, Table A.8 shows that our findings hold when using alternative measures of parental responses. To facilitate comparison across definitions, the first column of panel (a) replicates our main point estimate from Table 1. In this specification, our outcome variable is the extent to which parents help their children with homework, and we have standardized the survey answers by country and database. The remaining columns of the panel show that we still find significant compensatory responses if we instead normalize the survey responses to a 0–1 scale (second column), convert the possible survey responses to number of days per week (third column), or create a binary indicator for helping children every day or almost every day (fourth column).

## 6 Conclusions

This paper provides international evidence on whether parents and schools invest in children in ways that reinforce or compensate for academic performance. Using data on fourth and fifth graders in 34 countries, we exploit age-based school entry rules and show that children who are relatively young for their school cohort underperform relative to their older peers. Both parents and schools respond to children’s expected relative age in a compensatory way. For children who are expected to be young for grade, parents invest additional time; helping with schoolwork, practicing basic skills, and monitoring homework completion. Similarly, schools place children in smaller classes and provide remedial tutoring to a greater extent. Our preferred interpretation is that these behaviors reflect the compensatory responses to poorer academic performance.

Compensatory parental investments are a pervasive phenomenon, with significant effects in a wide variety of education systems around the world. We do not find differential responses by children’s gender or foreign background, although high-educated parents compensate to a

slightly larger extent than low-educated parents. Additionally, we provide suggestive evidence that parents' compensatory responses are stronger in less compensatory educational contexts. This aligns with previous studies finding that parental and school investments are substitutes.

We also shed some light on the mechanisms, using longitudinal data from the US to examine when compensatory responses emerge and whether they change across grade levels. We find that parental responses are fairly stable from first through fifth grade. Interestingly, however, parents do not compensate for worse academic performance around the time of kindergarten entry. One explanation may be that parents underestimate the magnitude or persistence of relative age effects on performance; in such case, providing better information could lead them to provide compensatory support at young ages—a time when such investments may have particularly beneficial effects.

Finally, schools do not appear to compensate for achievement disadvantages at kindergarten entry. School compensation then grows over grade levels, particularly for class size. These dynamics provide insight into why the effects of educational interventions may fade out over grade levels.

## References

- Agostinelli, Francesco and Matthew Wiswall**, “Estimating the Technology of Children’s Skill Formation,” *Journal of Political Economy*, 2025, 133 (3), 846–887.
- Angrist, Joshua D. and Victor Lavy**, “Using Maimonides’ Rule to Estimate the Effect of Class Size on Scholastic Achievement,” *The Quarterly Journal of Economics*, 05 1999, 114 (2), 533–575.
- Bailey, Drew H., Greg J. Duncan, Flávio Cunha, Barbara R. Foorman, and David S. Yeager**, “Persistence and Fade-Out of Educational-Intervention Effects: Mechanisms and Potential Solutions,” *Psychological Science in the Public Interest*, 2020, 21 (2), 55–97.
- Becker, Gary S.**, *A Treatise on the Family*, Harvard University Press, 1981.
- **and Nigel Tomes**, “Child Endowments and the Quantity and Quality of Children,” *Journal of Political Economy*, 1976, 84 (4, Part 2), S143–S162.
- Bedard, Kelly and Elizabeth Dhuey**, “The Persistence of Early Childhood Maturity: International Evidence of Long-Run Age Effects,” *The Quarterly Journal of Economics*, 11 2006, 121 (4), 1437–1472.
- Behrman, Jere R., Robert A. Pollak, and Paul Taubman**, “Parental Preferences and Provision for Progeny,” *Journal of Political Economy*, 1982, 90 (1), 52–73.
- Berne, Jordan S., Brian A. Jacob, Christina Weiland, and Katharine O. Strunk**, “The Impacts of Grade Retention Policy With Minimal Retention,” Working Paper 33764, National Bureau of Economic Research May 2025.
- Berniell, Inés and Ricardo Estrada**, “Poor Little Children: The Socioeconomic Gap in parental responses to School Disadvantage,” *Labour Economics*, 2020, 66, 101879.
- Björklund, Anders, Mikael Lindahl, and Erik Plug**, “The Origins of Intergenerational Associations: Lessons from Swedish Adoption Data,” *The Quarterly Journal of Economics*, 08 2006, 121 (3), 999–1028.
- Celhay, Pablo and Sebastian Gallegos**, “Early Skill Effects on Parental Beliefs, Investments and Children Long-Run Outcomes,” *Journal of Human Resources*, 2023, 58 (4).
- Chambers, Jay G.**, “An Analysis of Resource Allocation in Public School Districts,” *Public Finance Quarterly*, 1978, 6 (2), 131–160.
- Chang, Simon, Deborah A. Cobb-Clark, and Nicolás Salamanca**, “Parents’ Responses to Teacher Qualifications,” *Journal of Economic Behavior & Organization*, 2022, 197, 419–446.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff**, “Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood,” *American Economic Review*, September 2014, 104 (9), 2633–2679.
- **, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez**, “Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *The Quarterly Journal of Economics*, 09 2014, 129 (4), 1553–1623.

- Cullen, Julie Berry, Brian A Jacob, and Steven Levitt**, “The Effect of School Choice on Participants: Evidence from Randomized Lotteries,” *Econometrica*, September 2006, 74 (5), 1191–1230.
- Cunha, Flavio and James Heckman**, “The Technology of Skill Formation,” *American Economic Review*, 2007, 97 (2), 31–47.
- Currie, Janet and Duncan Thomas**, “School Quality and the Longer-Term Effects of Head Start,” *The Journal of Human Resources*, 2000, 35 (4), 755–774.
- Das, Jishnu, Stefan Dercon, James Habyarimana, Pramila Krishnan, Karthik Muralidharan, and Venkatesh Sundararaman**, “School Inputs, Household Substitution, and Test Scores,” *American Economic Journal: Applied Economics*, January 2013, 5 (2), 29–57.
- Dhuey, Elizabeth and Stephen Lipscomb**, “Disabled or young? Relative age and special education diagnoses in schools,” *Economics of Education Review*, 2010, 29 (5), 857–872.
- , **David Figlio, Krzysztof Karbownik, and Jeffrey Roth**, “School Starting Age and Cognitive Development,” *Journal of Policy Analysis and Management*, 2019, 38 (3), 538–578.
- Dizon-Ross, Rebecca**, “Parents’ Beliefs about Their Children’s Academic Ability: Implications for Educational Investments,” *American Economic Review*, August 2019, 109 (8), 2728–2765.
- Elder, Todd E. and Darren H. Lubotsky**, “Kindergarten Entrance Age and Children’s Achievement: Impacts of State Policies, Family Background, and Peers,” *The Journal of Human Resources*, 2009, 44 (3), 641–683.
- Fan, Wei and Catherine Porter**, “Reinforcement or Compensation? Parental Responses to Children’s Revealed Human Capital Levels,” *Journal of Population Economics*, January 2020, 33 (1), 233–270.
- Figlio, David and Umut Özek**, “The Unintended Consequences of Test-Based Remediation,” *American Economic Journal: Applied Economics*, January 2024, 16 (1), 60–89.
- Fredriksson, Peter and Björn Öckert**, “Is Early Learning Really More Productive? The Effect of School Starting Age on School and Labor Market Performance,” IZA Discussion Papers 1659, Institute of Labor Economics (IZA) July 2005.
- , – , and **Hessel Oosterbeek**, “Long-Term Effects of Class Size,” *The Quarterly Journal of Economics*, 11 2013, 128 (1), 249–285.
- , – , and – , “Parental Responses to Public Investments in Children: Evidence from a Maximum Class Size Rule,” *Journal of Human Resources*, October 2016, 51 (4), 832–868.
- Frijters, Paul, David W. Johnston, Manisha Shah, and Michael A. Shields**, “Intrahousehold Resource Allocation: Do Parents Reduce or Reinforce Child Ability Gaps?,” *Demography*, 2013, 50 (6), 2187–2208.
- Grätz, Michael and Florencia Torche**, “Compensation or Reinforcement? The Stratification of Parental Responses to Children’s Early Ability,” *Demography*, 2016, 53 (6), 1883–1904.

- Grönqvist, Erik, Björn Öckert, and Jonas Vlachos**, “The Intergenerational Transmission of Cognitive and Noncognitive Abilities,” *Journal of Human Resources*, 2017, 52 (4), 887–918.
- Hanushek, Eric A, Babs Jacobs, Guido Schwerdt, Rolf van der Velden, Stan Vermeulen, and Simon Wiederhold**, “The Intergenerational Transmission of Cognitive Skills: An Investigation of the Causal Impact of Families on Student Outcomes,” Working Paper 29450, National Bureau of Economic Research November 2021.
- Holmlund, Helena, Mikael Lindahl, and Erik Plug**, “The Causal Effect of Parents’ Schooling on Children’s Schooling: A Comparison of Estimation Methods,” *Journal of Economic Literature*, September 2011, 49 (3), 615–51.
- Jackson, C. Kirabo, Rucker C. Johnson, and Claudia Persico**, “The Effects of School Spending on Educational and Economic Outcomes: Evidence from School Finance Reforms,” *The Quarterly Journal of Economics*, 2016, 131 (1), 157–218.
- Jacob, Brian A., Lars Lefgren, and David P. Sims**, “The Persistence of Teacher-Induced Learning,” *The Journal of Human Resources*, 2010, 45 (4), 915–943.
- Johnson, Rucker C. and C. Kirabo Jackson**, “Reducing Inequality through Dynamic Complementarity: Evidence from Head Start and Public School Spending,” *American Economic Journal: Economic Policy*, November 2019, 11 (4), 310–349.
- Karbownik, Krzysztof and Umut Özek**, “Setting a Good Example? Examining Sibling Spillovers in Educational Achievement Using a Regression Discontinuity Design,” *Journal of Human Resources*, September 2023, 58 (5), 1567–1607.
- Krueger, Alan B.**, “Experimental Estimates of Education Production Functions,” *The Quarterly Journal of Economics*, 05 1999, 114 (2), 497–532.
- **and Diane M. Whitmore**, “The Effect of Attending a Small Class in the Early Grades on College-test Taking and Middle School Test Results: Evidence from Project Star,” *The Economic Journal*, 2001, 111 (468), 1–28.
- Landersø, Rasmus Kløve, Helena Skyt Nielsen, and Marianne Simonsen**, “Effects of School Starting Age on the Family,” *Journal of Human Resources*, 2020, 55 (4), 1258–1286.
- Lee, Valerie E. and Susanna Loeb**, “Where Do Head Start Attendees End Up? One Reason Why Preschool Effects Fade Out,” *Educational Evaluation and Policy Analysis*, 1995, 17 (1), 62–82.
- McEwan, Patrick J. and Joseph S. Shapiro**, “The Benefits of Delayed Primary School Enrollment: Discontinuity Estimates Using Exact Birth Dates,” *The Journal of Human Resources*, 2008, 43 (1), 1–29.
- Nicoletti, Cheti and Valentina Tonei**, “Do Parental Time Investments React to Changes in Child’s Skills and Health?,” *European Economic Review*, 2020, 127, 103491.
- Pop-Eleches, Cristian and Miguel Urquiola**, “Going to a Better School: Effects and Behavioral Responses,” *American Economic Review*, June 2013, 103 (4), 1289–1324.

- Rothstein, Jesse**, “Inequality of Educational Opportunity? Schools as Mediators of the Inter-generational Transmission of Income,” *Journal of Labor Economics*, January 2019, 37 (S1), S85–S123.
- Shapiro, Anna**, “Over Diagnosed or Over Looked? The Effect of Age at Time of School Entry on Students Receiving Special Education Services,” *Exceptional Children*, 2023, 89 (2), 161–177.
- Thomas, J. A.**, “Resource Allocation in School Districts and Classrooms,” *Journal of Education Finance*, 1980, 5 (3), 246–261.
- Todd, Petra E. and Kenneth I. Wolpin**, “On the Specification and Estimation of the Production Function for Cognitive Achievement,” *The Economic Journal*, 2003, 113 (485), F3–F33.
- Yurk Quadlin, Natasha**, “When Children Affect Parents: Children’s Academic Performance and Parental Investment,” *Social Science Research*, 2015, 52, 671–685.

# A. Main appendix

## A.1 Theoretical framework

Here we prove the claims of the main text. We also briefly discuss when schools exhibit compensatory behavior.

If parents have perfect foresight, we can think of parents as solving a two-stage problem where they first determine the optimal total amount of resources allocated to the child. Then, in a second stage, they decide on the optimal allocation of resources over time. The solution to the second problem is

$$\frac{p_1}{p_0} = \left( \frac{\mu_1}{\mu_0} \right)^{1/(1-\tau)}$$

Given the optimal allocation over time we have  $R_p = \delta(p_0 + p_1)$ , where  $\delta$  is a constant. The total amount of resources invested in the child is then determined by the first-order condition

$$\frac{\delta w}{w^p} \frac{\partial H}{\partial R_p} - MRS(c^p, y) = 0$$

For simplicity, we impose  $\delta w = w^p$  and assume  $h(R) = R^\rho$ . We can then write the first order condition as

$$\psi(R_p, H, s) \equiv \rho \kappa \frac{H}{R_p} - \frac{(wH)^\gamma}{c^p} = 0$$

where

$$\kappa = \frac{\mu_p R_p^\phi}{\mu_p R_p^\phi + \mu_s s^\phi}, \quad \kappa \in [0, 1], \quad (6)$$

$H = H(R_p, H, s)$ , and  $c^p = y^p - (w^p/\delta)R_p$  by the budget constraint.

## Compensatory behavior

By the implicit function theorem, we have

$$\frac{\partial R_p}{\partial x} = -\frac{\psi_x}{\psi_{R_p}}$$

where  $x = \{s, H\}$ . Since  $\psi_{R_p} < 0$ , the sign of this derivative is the same as the sign of  $\psi_x$ . We have

$$\text{sign} \frac{\partial R_p}{\partial H} = -\text{sign}\{\gamma - 1\}$$

and

$$\text{sign} \frac{\partial R_p}{\partial s} = -\text{sign}\{\phi + (\gamma - 1)\rho\}$$

Thus parents are compensatory— $\partial R_p/\partial H < 0$ —if and only if  $\gamma > 1$ . And if parents are

compensatory (i.e.,  $\gamma > 1$ ),  $\partial R_p/\partial s$  is unambiguously negative if  $\phi > 0$ , i.e., if private and public investments are substitutes. If so, we predict that parental investments fall in more compensatory school contexts.

## Parental information

If parents have imperfect information, the problem is no longer separable in the way described above.  $p_0$  is determined by expectations at  $t = 0$ , while  $p_1$  is determined after uncertainty has been revealed.

The parent-child trade-off—equation (1) in the main text—yields  $p_1 = p_1(p_0, a)$  where  $a$  is a shifter of the human capital production function.  $p_0$  is determined by the signal available as of time zero,  $a_0$ . We think of the signal as containing random measurement error. Thus,

$$E(a|a_0) = (1 - \alpha)E(a) + \alpha a_0$$

where  $0 \leq \alpha \leq 1$ . The condition determining the optimal allocation over time still holds in expectation:

$$p_0 = \left( \frac{\mu_0}{\mu_1} \right)^{1/(1-\tau)} E(p_1|a_0) = \tilde{\mu} E(p_1|a_0) \quad (7)$$

Suppose  $da < 0$ , which changes  $p_1$  by

$$dp_1 = \frac{\partial p_1}{\partial p_0} dp_0 + \frac{\partial p_1}{\partial a} da$$

where  $\partial p_1/\partial p_0$  is the slope of parent-child trade-off in  $(p_1, p_0)$ -space and  $\partial p_1/\partial a$  is the change in  $p_1$ , holding  $p_0$  constant.

How much  $p_0$  changes depends on the precision of the signal, however. Linearizing  $E(p_1|a_0)$ , we have

$$dp_0 = \tilde{\mu} \frac{\partial p_1}{\partial a} \frac{\partial E(a|a_0)}{\partial a_0} da_0$$

Now, consider a case where  $da = da_0 < 0$ . As long as the signal contains some information,  $dp_0 > 0$ . However, the increase is less than optimal, which will imply additional compensation in period 1. But the additional compensation is not sufficient to undo the fact that  $dp_0 < dp_0^*$ , where  $dp_0^*$  denotes the optimal adjustment given perfect foresight (i.e.  $\alpha = 1$ ). It is straightforward to show that

$$dp^* - dp = \left( 1 + \frac{\partial p_1}{\partial p_0} \right) \tilde{\mu} (1 - \alpha) dp_1^* \geq 0$$

where  $dp^* = (dp_1^* + dp_0^*)$  denotes optimal investments and  $dp = (dp_1 + dp_0)$  denotes the constrained optimum, given that  $p_0$  was suboptimal. The conclusion follows, since  $\alpha \leq 1$ .

## School behavior

Let us, finally, briefly describe the behavior of schools. Schools care about aggregate human capital as well as the distribution across the children in the school, and they take parental behavior as given. We take the school objective function to be

$$U^s = \left( \sum H_i^\lambda \right)^{1/\lambda}$$

where the CES parameter  $\lambda$  indexes how much the school cares about efficiency relative to inequality. If  $\lambda = 1$ , the school only cares about aggregate human capital; if  $\lambda \rightarrow -\infty$ , it cares only about equity.<sup>1</sup>

The school operates under a given budget constraint and thus only decides how to allocate resources across the student population:  $\sum s_i = \bar{s}$ . The optimal allocation of resources across students is governed by

$$\frac{s_i}{s_j} \frac{1 - \kappa_j}{1 - \kappa_i} = \left( \frac{H_i}{H_j} \right)^\lambda$$

where  $\kappa$  is defined in equation 6.

Suppose there is an innovation in  $H_i$ . Whether this causes a reallocation of resources away from the child, depends on the relative strength of efficiency and equity concerns in the school objective function – i.e., on the parameter  $\lambda$ . An innovation in  $H_i$  has two effects. First, it implies that it is more efficient to invest in child  $i$ . Second, it creates an incentive for schools to reallocate resources to other children, since the school values equality. When  $\lambda = 0$ , these two motives exactly balance, and the allocation of resources does not change in response to a change in  $H_i$ . Thus

$$\text{sign} \left\{ \frac{\partial s}{\partial H_i} \right\} = \text{sign} \{ \lambda \}$$

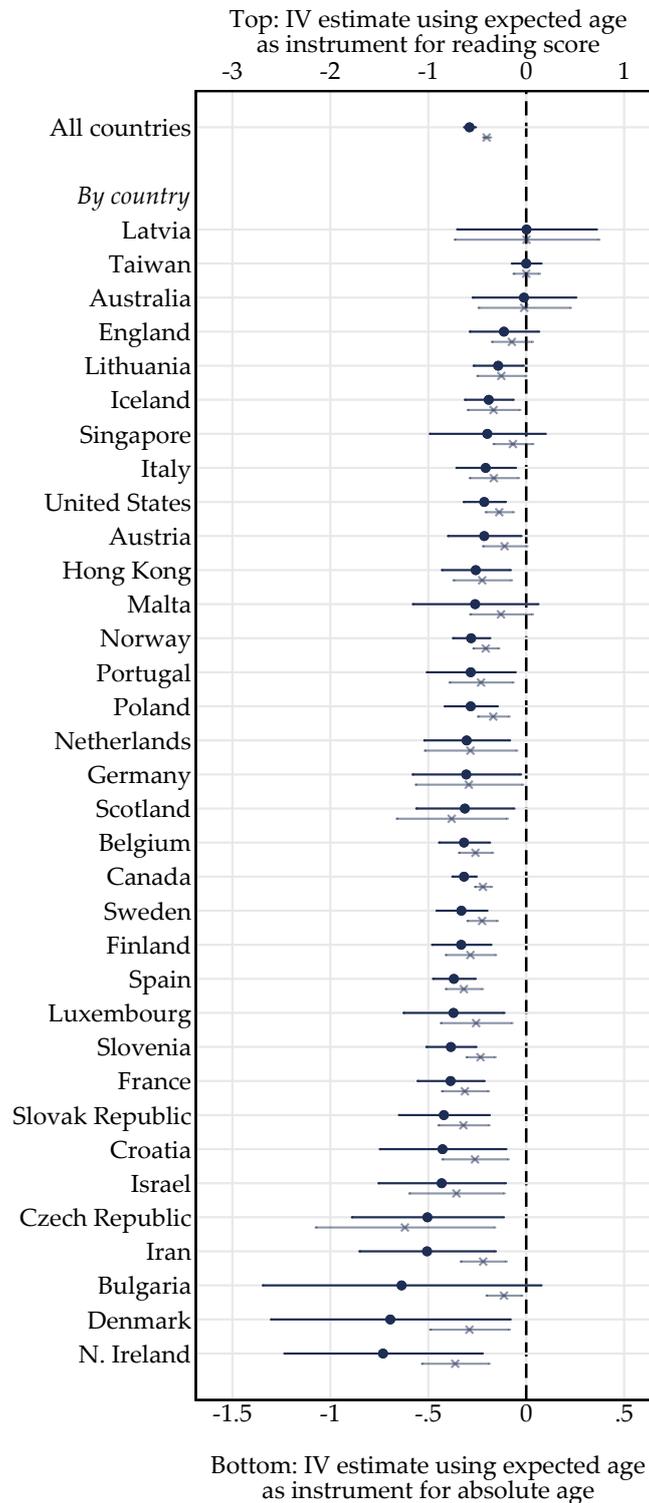
Schools must care sufficiently about equality for there to be a compensatory reduction in resources in response to an innovation in  $H_i$ .

---

<sup>1</sup>We do not explicitly model where these preferences come from, but we think of them as being generated from an initial state where parents agree on the objective function under the veil of ignorance.

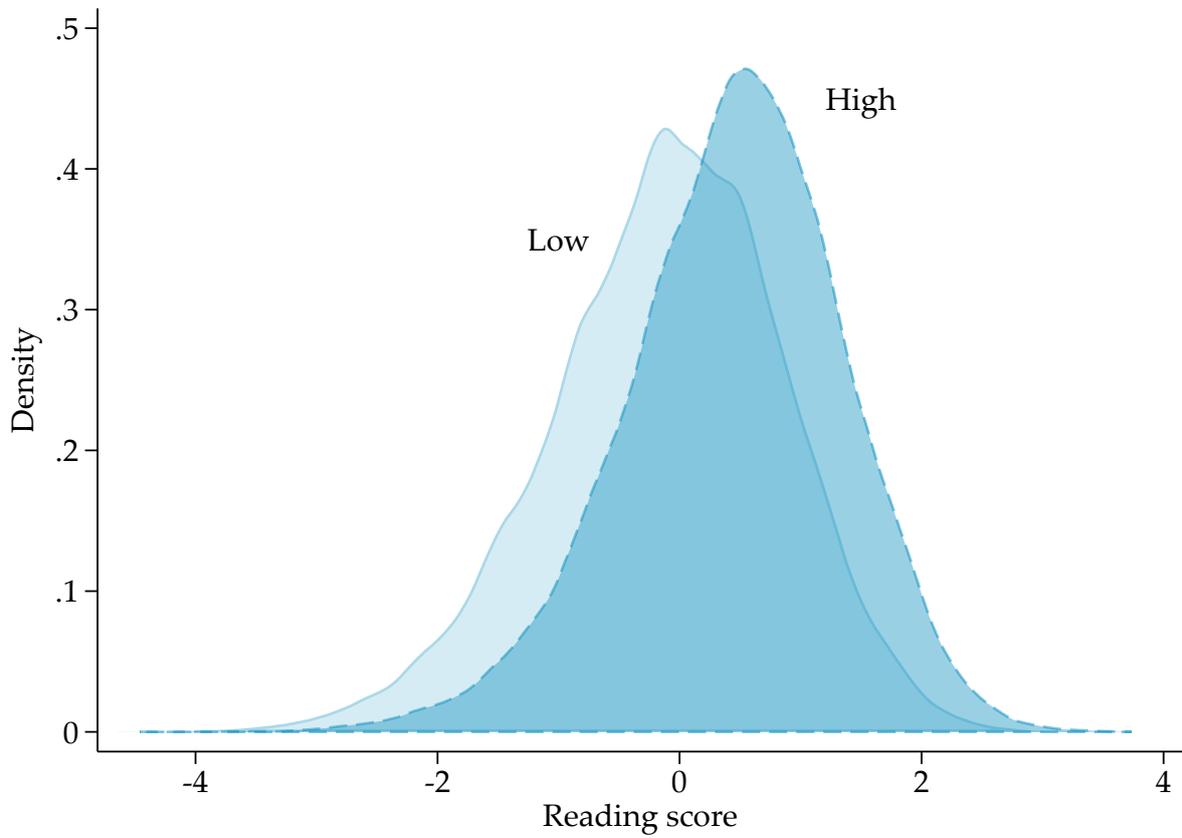


Figure A.2: Comparison of IV estimates for parental help with homework when using different scaling factors



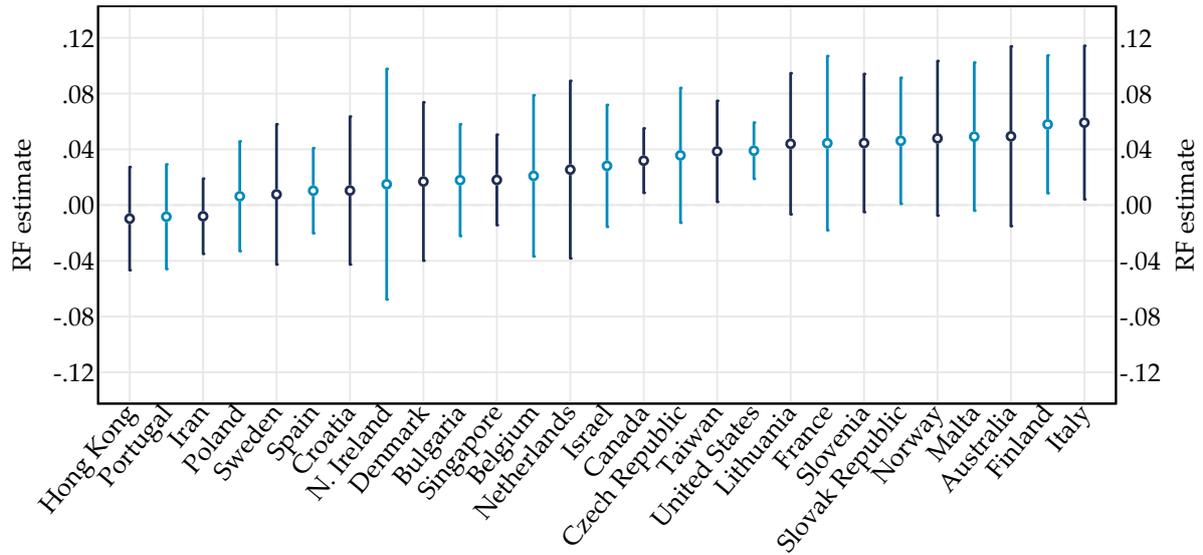
Notes: The first set of estimates corresponds to our main pooled sample; the remaining sets are country-specific. In each set, the top estimate, denoted with a dark circle, shows the IV coefficient and 95% confidence interval using students' expected relative age rank as an instrument for their standardized reading score. The bottom estimate, denoted with cross markers, shows the IV coefficient and 95% confidence interval using actual age rank as an instrument for expected relative age. The data used in the estimations are from PIRLS 2006, PIRLS 2011, and the fifth-grade waves of ECLS-K:1999 and ECLS-K:2011.

Figure A.3: Reading score distribution by parental education



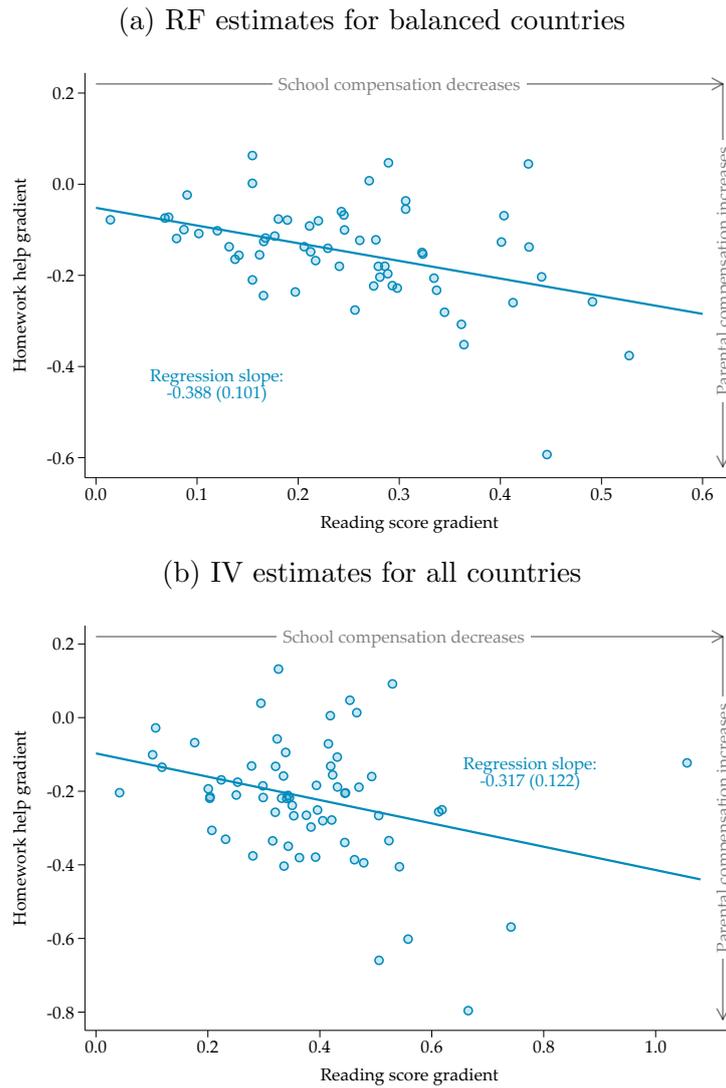
*Notes:* This figure shows the distribution of standardized reading scores for children with low-educated parents (lighter shading) and high-educated parents (darker shading). We define a child as having highly educated parents if either the mother or father holds a college degree. On average, children with high-educated parents score 0.56 standard deviations higher than children with low-educated parents.

Figure A.4: Parents' educational expectations across countries



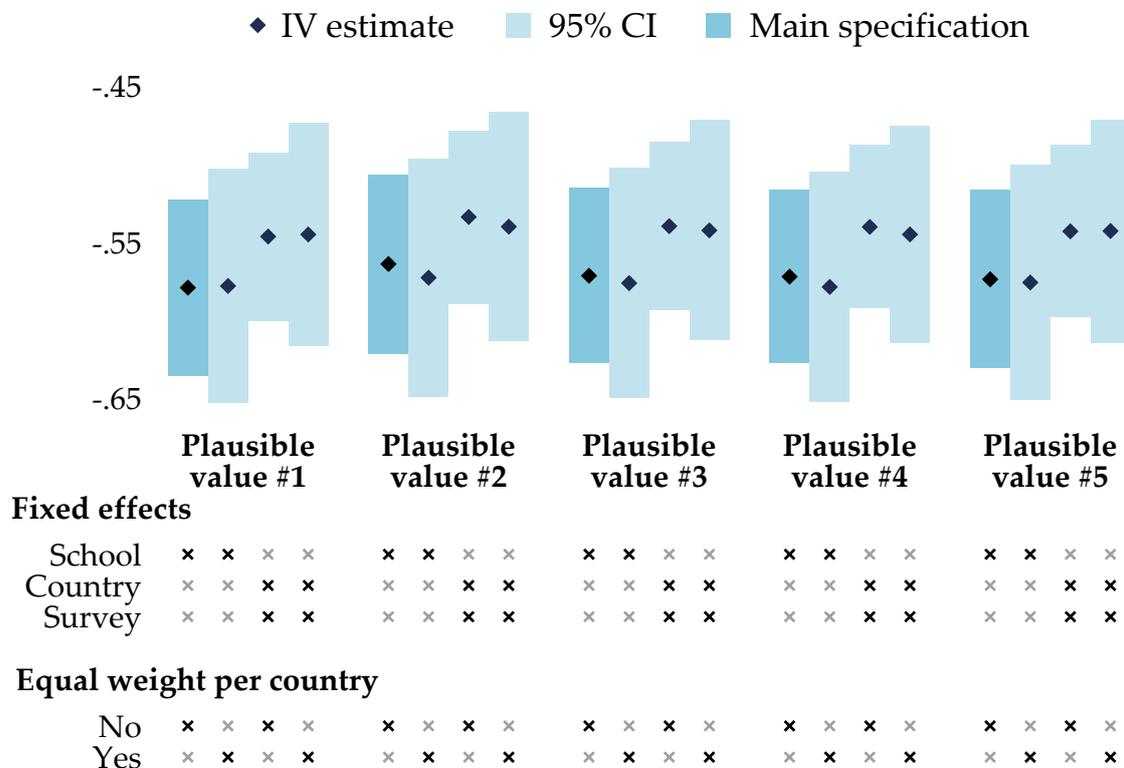
*Notes:* This figure shows the reduced-form relationship between children's expected relative age rank and parents' educational expectations. The outcome variable is an indicator equal to one if the parent expects that their child will complete a college degree. Data comes from PIRLS 2011 and the fifth-grade waves of ECLS-K:1999 and 2011. All regressions include school fixed effects. Standard errors are clustered by expected date of birth.

Figure A.5: Robustness checks for parental responses in different contexts



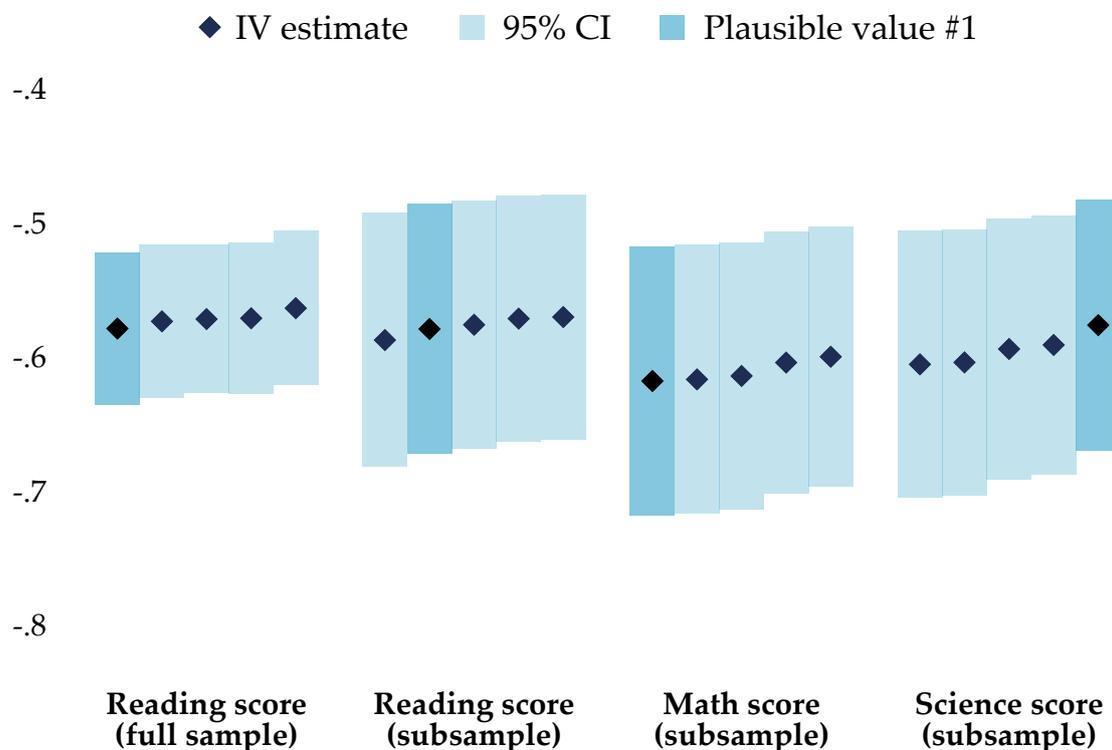
*Notes:* Panel (a) repeats the analysis shown in Figure 10 after dropping the three nations that fail the balance test. The reported slope comes from a regression of the regional homework help gradient on the regional reading gradient. There are 62 observations (two per country). Each observation is weighted by the inverse standard error of the reading score gradient. Panel (b) repeats the analysis shown in Figure 10 using an IV specification in which children's expected age rank is an instrument for their actual age rank. Each dot shows the IV estimates for one region, defined as a combination of country x school urbanicity. The reported slope comes from a regression of the regional homework help gradient on the regional reading gradient. There are 67 observations. Each observation is weighted by the inverse standard error of the reading score gradient.

Figure A.6: Robustness checks for parental homework help (std.) using different specifications and plausible values



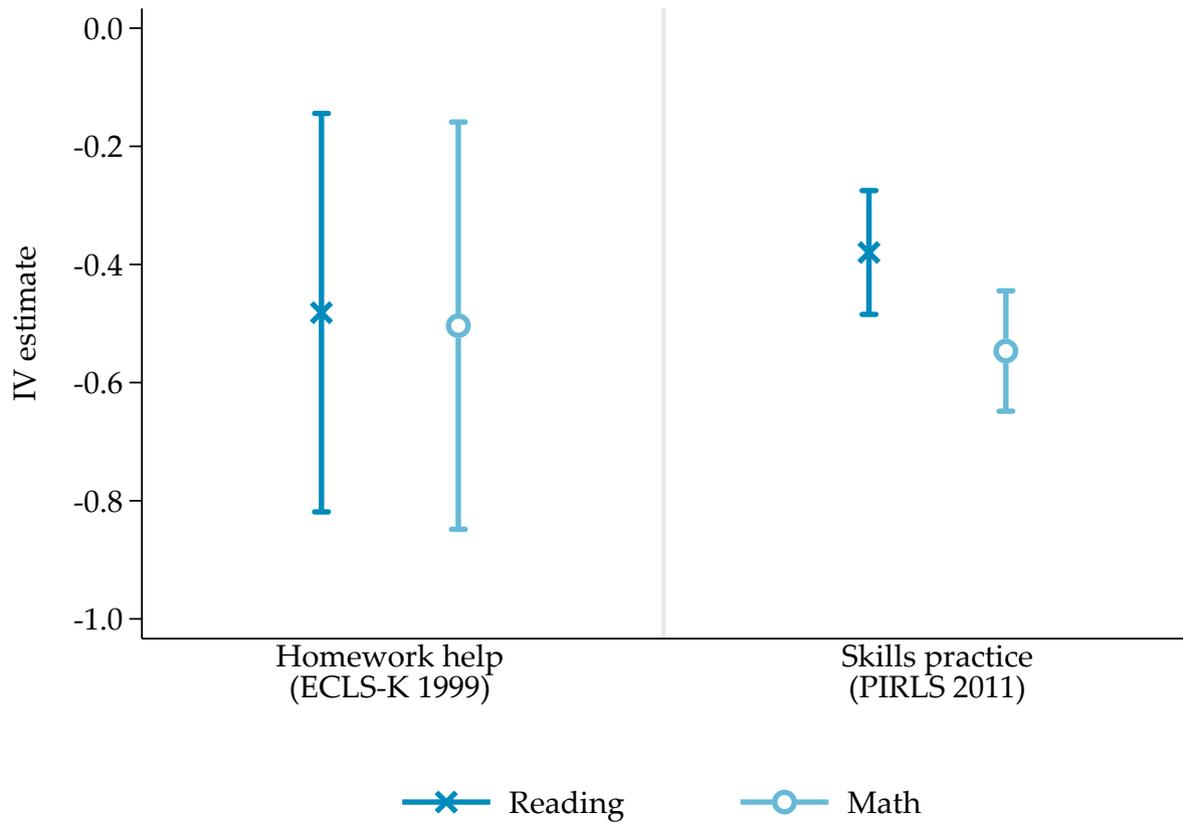
Notes: The figure plots the IV estimates and corresponding 95% confidence intervals from various regression specifications in which we use students' expected age rank as an instrument for one of the five plausible values of the PIRLS reading score. In the ECLS-K data, there is only one value for the reading score, and thus, this value is used for children in the US across all specifications. Standard errors are clustered by expected date of birth.

Figure A.7: Robustness checks for parental homework help (std.) using different test domains in the first stage



*Notes:* The figure plots the IV estimates and corresponding 95% confidence intervals from regressions in which we use students' expected age rank as an instrument for one of the five plausible values from the various subject tests in PIRLS, TIMSS, and ECLS-K. In the ECLS-K data, there is only one value for each subject, and thus, this value is used for children in the US across all specifications. The "full sample" refers to our main estimation sample. The "subsample" refers to the sample of students tested in all three subjects. Data from PIRLS 2006 is excluded from the subsample because TIMSS was not administered that year. Only a subset of data from PIRLS 2011 is included because not all countries participated in both TIMSS and PIRLS in 2011.

Figure A.8: Subject-specific parental responses (std.)



*Notes:* The figure plots the IV estimates and corresponding 95% confidence intervals from regressions in which we use students' expected age rank as an instrument for their reading score (darker bars) or math score (lighter bars). In the left panel, the outcome is the frequency of homework help in reading, language arts, or spelling (darker shading) or homework help in math (lighter shading). In the right panel, the outcome is either the frequency of reading practice (darker shading) or practice with math skills (lighter shading). We cannot use data from PIRLS 2006 or ECLS-K:2011 because there were no subject-specific questions about homework help and skills practice.

## A.3 Tables

Table A.1: Number of observations and survey response rates by country

Country	Obs.	Survey response rate	
		Parents	Teachers
All	287,721	0.875	0.944
Australia	5,059	0.536	0.681
Austria	9,737	0.941	0.988
Belgium	12,754	0.916	0.964
Bulgaria	9,111	0.964	0.979
Canada	35,733	0.838	0.923
Croatia	4,587	0.986	0.972
Czech Republic	4,553	0.967	0.996
Denmark	8,595	0.933	0.979
England	4,033	0.460	0.930
Finland	4,640	0.948	0.962
France	8,839	0.914	0.977
Germany	7,899	0.855	0.944
Hong Kong	8,585	0.943	0.986
Iceland	5,052	0.730	0.874
Iran	5,755	0.975	1.000
Israel	4,186	0.777	0.870
Italy	7,769	0.931	0.979
Latvia	4,160	0.935	0.977
Lithuania	9,361	0.954	0.995
Luxembourg	5,101	0.915	0.992
Malta	3,596	0.904	0.973
Netherlands	8,151	0.617	0.914
Northern Ireland	3,586	0.583	0.834
Norway	8,834	0.906	0.831
Poland	9,859	0.970	0.956
Portugal	4,085	0.946	0.954
Scotland	3,775	0.514	0.898
Singapore	6,367	0.970	0.979
Slovak Republic	10,992	0.966	0.984
Slovenia	9,848	0.948	0.992
Spain	12,631	0.824	0.952
Sweden	9,013	0.891	0.857
Taiwan	8,882	0.970	0.989
United States	22,593	0.841	0.935

*Notes:* The parental response rate is measured for the homework help question. The teacher response rate is measured for the class size question. For the US data, we report the number of observations and the response rates in the sub-sample of children who completed the fifth grade reading assessment.

Table A.2: Descriptive statistics for the different samples

	<i>International</i>	<i>United States data by grade level</i>						
	Grade 4–5 <sup>†</sup>	K, fall	K, spring	Grade 1	Grade 2 <sup>‡</sup>	Grade 3	Grade 4 <sup>‡</sup>	Grade 5
<b>Panel (a): Parental and school responses</b>								
Homework help (days/week)	2.77 (2.05)	n.a. n.a.	n.a. n.a.	3.46 (1.89)	3.44 (1.83)	3.30 (1.96)	2.78 (1.80)	2.32 (1.78)
Skills practice (days/week)	2.52 (1.90)	5.79 (1.70)	n.a. n.a.	5.42 (1.88)	4.87 (1.99)	5.18 (2.09)	3.76 (2.36)	3.27 (2.42)
Class size	23.77 (6.10)	20.30 (4.46)	20.35 (4.51)	20.93 (4.30)	21.47 (4.36)	21.69 (4.48)	22.90 (5.47)	23.04 (5.85)
Remedial tutoring	n.a.	n.a.	0.09	0.15	0.20	0.19	0.26	0.16
<b>Panel (b): Children’s background characteristics</b>								
Expected age at test	10.21 (0.53)	5.53 (0.32)	6.03 (0.32)	7.08 (0.32)	8.04 (0.31)	9.03 (0.32)	10.04 (0.31)	11.00 (0.32)
Actual age at test	10.29 (0.60)	5.59 (0.37)	6.09 (0.37)	7.14 (0.36)	8.10 (0.37)	9.09 (0.36)	10.10 (0.37)	11.06 (0.36)
Girl	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49
Foreign	0.16	0.13	0.15	0.17	0.22	0.18	0.23	0.19
High-educated parents	0.32	0.28	0.28	0.29	0.40	0.29	0.40	0.30

*Notes:* We convert the categorical responses for homework help and skills practice to days-per-week measures by using the midpoint of each category. See the data appendix for additional details on our variable definitions. Table A.3 shows the number of observations in each cell. <sup>†</sup>In the international sample, just under 90% of children are in fourth grade at the time of assessment. Several countries administer the PIRLS test in third or fifth grade, and we use fifth-grade data from the US due to data availability. <sup>‡</sup>There is no data from ECLS-K:1999 in second or fourth grade.

Table A.3: Number of observations in the different samples

	<i>International</i>	<i>United States data by grade level</i>						
	Grade 4–5 <sup>†</sup>	K, fall	K, spring	Grade 1	Grade 2 <sup>‡</sup>	Grade 3	Grade 4 <sup>‡</sup>	Grade 5
<b>Panel (a): Parental and school responses</b>								
Homework help	251,596	n.a.	n.a.	26,358	11,348	11,930	9,658	18,965
Skills practice	133,776	12,904	n.a.	14,602	11,341	12,291	9,734	8,980
Class size	271,690	30,479	31,768	27,701	12,411	23,217	10,935	21,090
Remedial tutoring	n.a.	n.a.	33,747	27,837	12,463	23,067	10,872	20,693
<b>Panel (b): Children’s background characteristics</b>								
Expected age at test	287,721	33,281	36,112	31,443	13,833	27,141	12,071	22,690
Actual age at test	287,721	33,281	36,112	31,443	13,833	27,141	12,071	22,690
Girl	287,721	33,248	36,082	31,422	13,815	27,125	12,057	22,678
Foreign	246,136	31,576	34,221	30,015	13,306	25,986	11,670	21,814
High-educated parents	246,507	31,403	34,112	29,926	13,135	25,925	11,537	21,756

*Notes:* Each cell of the table reports the number of observations used to compute the descriptive statistics in the corresponding cell in Table A.2. <sup>†</sup>In the international sample, just under 90% of children are in fourth grade at the time of assessment. Several countries administer the PIRLS test in third or fifth grade, and we use fifth-grade data from the US due to data availability. <sup>‡</sup>There is no data from ECLS-K:1999 in second or fourth grade.

Table A.4: Average school characteristics

Variable	Estimate	Std. Error	Average	Obs.
School enrollment (total)	1.029	2.342	514.661	267,225
Urban school	-0.008	0.006	0.356	269,545
Share low SES students in school	-0.001	0.002	0.241	259,914
Class size	0.034	0.081	23.779	278,321
Teacher's years of experience	0.025	0.089	17.290	278,693
Age when tested	0.031	0.004	10.294	287,721
Girl	0.001	0.001	0.493	287,721
Foreign background	-0.001	0.003	0.166	286,562
High-educated parent(s)	0.001	0.001	0.318	286,586
Parent(s) in a professional job	0.003	0.002	0.412	264,027
Reading score (std.)	0.024	0.004	0.001	287,721

*Notes:* Each row reports the results of a separate regression akin to the reduced-form model in equation (5), but omitting the school fixed effects and using average school-level characteristics as the dependent variable. We use pooled international data from PIRLS 2006, PIRLS 2011, and the fifth-grade waves of ECLS-K:1999 and ECLS-K:2011. All regressions include fixed effects for country and database. Standard errors are clustered by expected date of birth.

Table A.5: Robustness check – dropping countries that fail the balance test

	Homework help		Class size	
	All countries (1)	Balanced countries (2)	All countries (3)	Balanced countries (4)
IV estimate	-0.579 (0.029)	-0.580 (0.030)	0.218 (0.053)	0.260 (0.054)
Mean	0.009	0.010	23.772	23.153
Observations	251,596	232,799	271,690	252,534

*Notes:* Columns (1) and (3) report our main IV estimates. Columns (2) and (4) report the corresponding IV estimates after dropping the three nations (Croatia, Hong Kong, and Singapore) that fail the balance test in Figure 3b. All regressions include school fixed effects. Standard errors are shown in parentheses and clustered by expected date of birth.

Table A.6: Amount of homework assigned by the teacher

	Assigned any reading HW		Frequency (days/week)		Expected time (min/assignment)	
	(1)	(2)	(3)	(4)	(5)	(6)
Expected age rank	0.002 (0.001)	0.002 (0.001)	-0.007 (0.011)	0.006 (0.005)	-0.061 (0.120)	0.008 (0.045)
Mean	0.958		2.943		22.273	
Observations	268,768		258,300		271,452	
Fixed effects:						
Country	✓	-	✓	-	✓	-
Database	✓	-	✓	-	✓	-
School	-	✓	-	✓	-	✓

*Notes:* All regressions control for grade level and are estimated using pooled international data from PIRLS and ECLS-K. Standard errors are shown in parentheses and clustered by expected date of birth. The data come from class-level teacher questionnaires. Questions vary slightly across surveys. The sample size drops in columns (3)–(4) because the teacher survey in ECLS-K:1999 did not include a question on assigned days of homework per week.

Table A.7: Country-level estimates for parental help with homework

Country	Obs.	First stage (standardized)	Reduced form (standardized)	IV specification (standardized)	Mean days of help per week
All	251,596	0.239 (0.008)	-0.138 (0.007)	-0.579 (0.029)	2.774 (2.049)
Australia	2,711	0.251 (0.063)	-0.006 (0.067)	-0.023 (0.267)	2.648 (1.868)
Austria	9,163	0.125 (0.025)	-0.053 (0.025)	-0.429 (0.180)	2.309 (2.081)
Belgium	11,687	0.227 (0.029)	-0.144 (0.023)	-0.635 (0.130)	2.802 (2.106)
Bulgaria	8,785	0.065 (0.025)	-0.083 (0.034)	-1.272 (0.721)	2.899 (2.075)
Canada	29,954	0.287 (0.017)	-0.182 (0.017)	-0.635 (0.061)	3.058 (2.002)
Croatia	4,521	0.187 (0.049)	-0.160 (0.052)	-0.854 (0.326)	2.968 (1.969)
Czech Republic	4,405	0.163 (0.051)	-0.164 (0.060)	-1.009 (0.393)	3.498 (1.892)
Denmark	8,016	0.071 (0.031)	-0.098 (0.037)	-1.388 (0.620)	2.989 (1.927)
England	1,853	0.311 (0.051)	-0.071 (0.045)	-0.227 (0.159)	2.108 (1.888)
Finland	4,400	0.356 (0.053)	-0.236 (0.054)	-0.664 (0.152)	2.048 (1.781)
France	8,081	0.267 (0.033)	-0.206 (0.040)	-0.772 (0.172)	3.423 (2.004)
Germany	6,756	0.169 (0.032)	-0.104 (0.046)	-0.612 (0.250)	1.771 (1.944)
Hong Kong	8,099	0.223 (0.027)	-0.115 (0.038)	-0.515 (0.177)	2.515 (2.128)
Iceland	3,686	0.430 (0.051)	-0.164 (0.064)	-0.382 (0.119)	3.165 (2.084)
Iran	5,613	0.165 (0.037)	-0.167 (0.045)	-1.012 (0.352)	2.958 (2.029)
Israel	3,253	0.210 (0.050)	-0.182 (0.062)	-0.863 (0.329)	2.788 (1.968)
Italy	7,232	0.258 (0.039)	-0.107 (0.041)	-0.414 (0.153)	3.310 (2.031)
Latvia	3,889	0.186 (0.051)	0.001 (0.060)	0.003 (0.322)	1.876 (2.001)
Lithuania	8,926	0.237 (0.045)	-0.068 (0.037)	-0.287 (0.128)	2.508 (2.074)
Luxembourg	4,669	0.167 (0.047)	-0.124 (0.040)	-0.743 (0.232)	2.406 (2.081)
Malta	3,251	0.209 (0.058)	-0.109 (0.068)	-0.522 (0.324)	3.011 (2.083)
Netherlands	5,032	0.169 (0.041)	-0.103 (0.044)	-0.608 (0.221)	1.456 (1.710)
Northern Ireland	2,091	0.231 (0.071)	-0.337 (0.081)	-1.461 (0.513)	4.105 (1.604)
Norway	8,006	0.357 (0.033)	-0.201 (0.032)	-0.563 (0.095)	3.397 (1.928)
Poland	9,566	0.278 (0.026)	-0.158 (0.038)	-0.567 (0.136)	3.073 (2.104)
Portugal	3,864	0.270 (0.049)	-0.153 (0.054)	-0.566 (0.229)	3.477 (1.927)
Scotland	1,937	0.423 (0.091)	-0.265 (0.075)	-0.627 (0.226)	2.716 (1.957)
Singapore	6,177	0.152 (0.042)	-0.060 (0.045)	-0.397 (0.299)	3.078 (1.952)
Slovak Republic	10,619	0.165 (0.032)	-0.139 (0.031)	-0.841 (0.234)	2.574 (2.092)
Slovenia	9,336	0.249 (0.036)	-0.191 (0.030)	-0.769 (0.128)	2.246 (2.023)
Spain	10,410	0.349 (0.031)	-0.258 (0.038)	-0.739 (0.109)	3.685 (1.901)
Sweden	8,030	0.295 (0.037)	-0.195 (0.033)	-0.661 (0.130)	2.517 (1.911)
Taiwan	8,613	0.409 (0.040)	-0.000 (0.031)	-0.000 (0.075)	2.447 (2.009)
United States	18,965	0.235 (0.024)	-0.101 (0.026)	-0.428 (0.108)	2.321 (1.785)

*Notes:* The first-stage column presents estimates of  $\gamma$  from equation (4), where the outcome is student  $i$ 's reading score (standardized by country and wave). The reduced-form column presents estimates of  $\pi^p$  from equation (5), where the outcome is the frequency that student  $i$  received homework help from their parents (standardized by country and survey wave). The IV-specification column reports the reduced form scaled by the first stage ( $\beta^p = \pi^p/\gamma$ ). US estimates are based on data from the fifth-grade waves of ECLS-K:1999 and 2011. All other estimates are based on data from PIRLS 2006 and 2011. The regressions include school fixed effects. Standard errors are reported in parentheses and clustered by expected date of birth. The final column reports the average number of days that parents help with homework per week, with standard deviations shown in parentheses.

Table A.8: Different measures of parental inputs (IV estimates)

	Standardized by country	Normalized (0–1 scale)	Number of days/week	(Almost) every day
<b>Panel (a): Frequency of help with homework</b>				
Reading score (std.)	–0.579 (0.029)	–0.172 (0.008)	–1.037 (0.052)	–0.219 (0.013)
Mean	0.009	0.698	2.774	0.432
Observations	251,596	251,596	251,596	251,596
<b>Panel (b): Frequency of skills practice</b>				
Reading score (std.)	–0.503 (0.041)	–0.147 (0.012)	–0.731 (0.073)	–0.105 (0.018)
Mean	0.008	0.650	2.520	0.279
Observations	133,776	133,776	133,776	133,776
<b>Panel (c): Frequency of homework monitoring</b>				
Reading score (std.)	–0.373 (0.042)	–0.088 (0.010)	–0.544 (0.072)	–0.121 (0.018)
Mean	0.007	0.876	4.007	0.746
Observations	134,344	134,344	134,344	134,344
<b>Panel (d): Frequency of discussion about schoolwork</b>				
Reading score (std.)	–0.126 (0.033)	–0.026 (0.008)	–0.198 (0.053)	–0.049 (0.013)
Mean	0.013	0.810	3.439	0.599
Observations	232,861	232,861	232,861	232,861
<b>Panel (e): Frequency of discussion about reading</b>				
Reading score (std.)	–0.116 (0.032)	–0.031 (0.009)	–0.153 (0.056)	–0.027 (0.014)
Mean	0.011	0.709	2.556	0.368
Observations	232,754	232,754	232,754	232,754

*Notes:* When possible, the robustness checks are estimated using pooled international data from multiple waves of PIRLS and ECLS-K. Due to data availability, panels (b) and (c) include data only from the 2011 studies. In column (1), parents' responses are standardized by database, country, and grade level. In column (2), the raw data is normalized to a 0–1 scale. In column (3), categorical responses are converted to days per week using the midpoint of the bin. In column (4), the outcome is a binary indicator equal to one if the parent reports doing the activity (almost) every day. The survey questions vary slightly across databases; see the data appendix for details.

## B. Data appendix

Our study relies on public-use data from the International Association for the Evaluation of Educational Achievement (IEA) and the United States' National Center for Education Statistics (NCES).<sup>1</sup> In this data appendix, we provide a comprehensive description of the databases, our sample restrictions, and the variables used in our analysis.

### B.1 Data sources

Data for all countries except the United States come from PIRLS, the Progress In International Reading Literacy Study. PIRLS is an internationally standardized reading assessment targeted at students who have completed four years of primary education and who are, on average, around 10 years old.<sup>2</sup> The IEA has administered the assessment every five years since 2001. Our analysis focuses on the 2006 and 2011 waves due to data limitations in the 2001, 2016, and 2021 waves.<sup>3</sup>

Although the US participated in PIRLS 2006 and 2011, the data available from the US are not suitable for our study. For example, we cannot analyze parental responses because the US opted not to administer the Home Questionnaire to the parents of participating students. To include the US in our analysis, we use another data source that contains comparable information on children and their parents, teachers, and schools.

The US data come from the Early Childhood Longitudinal Studies (ECLS) program, which is run by the NCES. The program includes two longitudinal studies, each of which tracked a nationally representative sample of students from kindergarten age through the end of primary school: ECLS-K:1999 (kindergarten cohort 1998/99) and ECLS-K:2011 (kindergarten cohort 2010/11).

### B.2 School-starting-age rules

Our empirical approach relies on school-starting-age (SSA) rules that stipulate the age at which children are old enough to begin formal schooling. These rules usually specify an exact cutoff date by which children must reach a certain age in order to enter primary school. In most countries, legislation on school-starting age is set at the national level, and the same cutoff date applies to all children in the country. There are, however, some countries where regulations

---

<sup>1</sup>The data is publicly available through the [IEA Data Repository](#) and [Early Childhood Longitudinal Studies \(ECLS\) Program – Data Products](#).

<sup>2</sup>In a few countries, students are assessed in their fifth year of schooling to ensure a minimum average age of 9.5 years old. This applies to students in England, Malta, New Zealand, and Trinidad and Tobago.

<sup>3</sup>In the 2001 wave, the Home Questionnaire did not ask about the extent to which parents help their children with homework or schoolwork. Thus, we cannot study our main parental response of interest using the data. Additionally, the public-use files for the 2016 and 2021 waves lack information on children's date of birth and the date of their reading assessment. This omission precludes us from deriving students' expected relative age, which is a crucial element of our identification strategy.

are set at the local level, and the cutoff date can vary across provinces, states, or even school districts.

In this section, we discuss the four countries in our sample that have region-specific SSA rules and explain how we determine the relevant rule for children in those countries.<sup>4</sup> We list the national SSA rules for the other countries in our sample in Table B.1. All information was obtained from questionnaires completed by PIRLS National Research Coordinators. If the rule reported by the Coordinator was ambiguous or not verifiable in the data, we excluded the region from our sample. Table B.2 summarizes our exclusion criteria and the regions dropped under each criterion.

**Australia** Each Australian state and territory sets its own policy on the age at which children begin primary school. Table B.3 provides an overview of the rules. In the PIRLS data, we identify the state or territory where children attend school – and thus the relevant SSA rule – through the explicit stratification variable. The strata codes are provided in the PIRLS 2011 User Guide (see Supplement 4, Exhibit S4.1 in Foy and Drucker, eds, 2013).

**Canada** The thirteen provinces and territories in Canada manage their own school system and set their own SSA rules. Table B.4 lists the rules for the nine provinces that participated in PIRLS 2006 and/or 2011. Most participating provinces stipulate that children must enter primary school the calendar year they turn six, implying a cutoff date of January 1. Only three of the participating provinces deviate from this rule: Alberta, Nova Scotia, and Quebec. Thus, to determine the relevant SSA rule for each child, we need to identify which children reside in these three provinces. All other children can be assigned a January cutoff.

In the 2006 wave, it is straightforward to identify where children attend school because each province was a benchmark participant. Consequently, the data were stored in separate province-level databases instead of a single country-level database. The filenames include a three-character abbreviation indicating which province the data come from.<sup>5</sup> We use these abbreviations to determine where children attended school at the time of the reading test and derive their expected date of birth according to the relevant SSA rule in that region. We exclude children in Alberta from our analysis because local school authorities are permitted to establish their own cutoff dates.

In 2011, only two of the three provinces that we need to identify – Alberta and Quebec – participated as benchmarking participants. As before, the data for benchmarking participants were stored in separate province-level files, enabling easy identification of the children who reside there. By contrast, the data from Nova Scotia were included with the data from the other non-

---

<sup>4</sup>In these cases, there is some measurement error in the rule that we assign to students because we cannot be certain where they lived when they entered primary school. For instance, some students in the PIRLS data may have moved between regions with different SSA rules by the time we observe them in fourth grade.

<sup>5</sup>The abbreviations for the benchmarking provinces are CAB (Alberta), CBC (British Columbia), CNS (Nova Scotia), COT (Ontario), and CQU (Quebec).

benchmarking participants in a set of country-level files. The files lack an indicator variable specifying where the data was collected, so we rely on a data-driven approach to identify which schools – and thus which children – are located in Nova Scotia.

Three pieces of information guide our approach for identifying schools in Nova Scotia: (i) suggestive evidence from the benchmarking provinces that school identifiers are numbered sequentially by region; (ii) the fact that children in Nova Scotia should be born between October 1995 and September 1996, assuming full compliance with the SSA rule; and (iii) official statistics reporting that 203 schools and 4,388 children participated in Nova Scotia (see Tables III.2 and III.3 in [Labrecque et al., eds, 2012](#)). Guided by this information, we study the date-of-birth distribution at the school level and find a sequence of 203 school identifiers in which the observed distribution complies with an October rather than January cutoff. In line with the official statistics, we observe 4,388 children in these schools.

**Germany** All children in Germany are required to attend school from age six onwards. The sixteen federal states, known as *Länder*, have the authority to establish distinct cutoff dates by which children must turn six to enroll in school. Historically, there was a uniform cutoff date of July 1 across the entire country. Between 2003 and 2011, half of the states gradually implemented later cutoff dates ([Schwandt and Wuppermann, 2016](#)). Due to the lack of regional identifiers in the PIRLS data, we cannot determine the relevant SSA cutoff date for children who entered first grade after the 2002/03 school year. We therefore restrict our analysis to the 2006 wave, as fourth graders in spring 2006 were the last cohort to face the uniform cutoff.

**United States** The age requirement to enter kindergarten varies across US states, and in a few instances, across local education authorities within a state. Ideally, our data would include information about children’s place of residence when they entered kindergarten. This would allow us to assign their expected date of birth according to the region’s SSA policy (see, e.g., [Elder and Lubotsky, 2009](#)). However, geographic identifiers are suppressed in the public-use versions of ECLS-K:1999 and ECLS-K:2011. We therefore use a data-driven approach to determine the applicable SSA rule for each child.

Our algorithm assigns children an SSA rule by analyzing the birth-date distribution in the school where they attended kindergarten. We compare the observed distribution to the expected distributions under various state-mandated cutoffs. Then, we select the cutoff that best aligns the observed distribution with the expected one. The procedure is detailed below.

1. We create an abbreviated set of state-level SSA rules that were in place when children in the sample entered kindergarten. The abbreviated set combines cutoff dates that are too close in time to distinguish using the observed birth-date distribution. For instance, cutoffs on September 30 and October 1 are grouped together; see [Table B.5](#) for the complete list.

2. We assign all children to a base-year school, typically the one attended during the fall term of kindergarten. For kindergarteners who entered the study after the fall term, we use the spring-term school as the base-year school.
3. We select a subsample of first-time kindergarteners, excluding children whose parents reported that they did not comply with the school’s SSA guidelines.
4. We iterate over all SSA rules in the abbreviated set and calculate the share of the subsample in each base-year school whose birth dates comply with the rule.
5. We assign children the SSA rule that maximizes the share of compliers in their base-year school. If multiple cutoff dates maximize the share of compliers, we select the cutoff date that is most common in the national distribution.<sup>6</sup>

The ECLS-K:1999 data also include information on SSA rules from the School Administrator Questionnaire.<sup>7</sup> In the first wave of the survey – i.e., the start of kindergarten – school administrators were asked, “*By what date did a child need to turn five to enter kindergarten for this school year, 1998–1999?*”. For consistency in variable definitions across studies, we do not use this data for our main analysis. We have, however, created a “reported cutoff” based on the survey responses and verified that our estimates are insensitive to the use of reported cutoffs rather than data-driven cutoffs (see Figure B.1).

## B.3 Variable definitions

This section describes how we construct several key variables used in our analysis. All variables are derived from interviews and survey data. Thus, we also provide details on the questions posed to respondents, as well as the possible answers they could give. The ECLS-K interview and survey items are reproduced from official documentation provided by the NCES.<sup>8</sup> We have copied the PIRLS survey items from official documentation provided by the IEA.<sup>9</sup>

### B.3.1 Date of birth

**ECLS-K** We observe children’s exact date of birth in ECLS-K:1999, but in ECLS-K:2011, we observe only the month of birth, as well as a binned variable for birth year. We can, however, approximate the exact date of birth for children in ECLS-K:2011 using information on age in

---

<sup>6</sup>For ECLS-K:1999, the tie-breaking order is as follows: September 1, December 1, October 1, January 1, October 15, June 1, August 1, September 15, and August 15. For ECLS-K:2011, the tie-breaking order is as follows: September 1, December 1, October 1, August 15, August 1, January 1, September 15, and October 15.

<sup>7</sup>The answers to these questions are suppressed in the public-use files from ECLS-K:2011.

<sup>8</sup>Source: National Center for Education Statistics. Early Childhood Longitudinal Studies Program (ECLS) — Instruments and Assessments.

<sup>9</sup>Source: International Association for the Evaluation of Educational Achievement, PIRLS 2006 Contextual Questionnaires and PIRLS 2011 Contextual Questionnaires.

months at school start, age in months at the time of various assessments, and the date of these events. As a validation exercise, we implemented the same approach for the ECLS-K:1999 data and compared children’s derived date of birth with their actual date of birth. Figure B.2 shows considerable overlap in the actual and derived distributions. The two dates are exact matches for 89% of the sample, and for the cases that do not match, the absolute difference is only 1.89 days on average. The results of this validation exercise suggest that there is minimal measurement error in the derived date of birth for children in ECLS-K:2011.

**PIRLS** We observe children’s month and year of birth in both PIRLS waves, but only the exact day of birth for the 2011 wave. We impute the day of birth as the 15th of the month for all missing cases. This adjustment affected only 4% of observations in PIRLS 2011 (0.7% outside of Austria, where the day of birth is always missing).

### B.3.2 Expected relative age

To derive children’s expected relative age, we use two pieces of information: (i) the SSA rule in their country or region, and (ii) their date of birth.

First, we define children’s “expected birthdate” by adjusting their actual year of birth such that the resulting birthdate complies with the SSA rule in their country or region. For example, consider a country that requires children to turn six before September 1, 2007 in order to enter primary school in the 2007/2008 school year. If all children comply with the SSA rule, and children cannot skip or repeat a grade level, then everyone in the grade would be born between September 1, 2000 and August 31, 2001. In practice, however, some children enter primary school early or late, and some may be held back or promoted. The expected birthdate for these non-compliers is found by setting the expected birth year to 2000 for everyone born from September through December and to 2001 for everyone born from January through August.

Next, we rank children based on their expected birthdate, separately by country/region and cohort. The ranking is normalized such that values of zero and one correspond to children who are expected to be youngest and oldest for their grade, respectively.

Note that children’s actual and expected birthdate—and thus their actual and expected age rank—are highly correlated ( $\rho \approx 0.7$  in our pooled international data). The correlation is stronger in regions with stricter school entry rules and limited grade retention or advancement.

### B.3.3 Student performance

We measure student performance using scores on standardized reading assessments that were administered as part of the ECLS-K and PIRLS studies. Children completed the assessments at school, but the results were solely for research purposes and did not affect their grades or academic progression. Because the results were never reported to participants of the study,

schools and parents could not adjust their investments in children as a direct reaction to test performance.

In the main manuscript, we provide two pieces of evidence that students' standardized reading score is a reasonable proxy for their overall school performance. First, in panel (a) of Figure 2, we show that the expected age gradient is similar across subjects using data from a subsample of children who completed standardized tests in multiple subjects around the same time.<sup>10</sup> Second, in panel (b) of Figure 2, we use ECLS-K data from the US to show that the expected relative age gradient is similar for other performance measures that could be more salient to parents and schools, such as teacher evaluations of children's reading proficiency<sup>11</sup> and teacher evaluations of children's socio-emotional skills.<sup>12</sup>

**ECLS-K** In each wave, the data include an Item Response Theory (IRT) theta score for children who participated in the reading assessment. The scores are an estimate of children's underlying reading skills, determined by their performance on the specific set of test items that they were administered. The higher the value, the higher the child's estimated ability. Although the children answered different sets of questions depending on their demonstrated ability,<sup>13</sup> the scores are adjusted to reflect a child's latent ability and are unaffected by the difficulty level of the questions that they answered. For our main analysis, we use the IRT theta scores in reading, standardized by database and grade level. Due to an error in the original ECLS-K:1999 data release, we use the corrected set of theta values provided by the NCES.

**PIRLS** To minimize the burden of the test, a limited number of assessment items were administered to each child. PIRLS uses IRT and multiple imputation techniques to derive estimates of what student performance on the assessment as a whole would have been, had the student completed the entire test. Five plausible values for students' overall reading score are included in the data. Our main analysis uses the first plausible value. We standardize the score by database, country, and—when relevant—grade level,<sup>14</sup> applying PIRLS sampling weights for representativeness. We also present robustness checks showing that the estimates are not sensitive to using

---

<sup>10</sup>The ECLS-K studies included standardized assessments for several subjects in every survey wave. Additionally, in 2011, some countries administered the Trends in International Mathematics and Science Study (TIMSS) to the same sample of children participating in PIRLS. We do not use the TIMSS data in our main analysis due to the lack of equivalent data for children in PIRLS 2006 and for some participating countries in PIRLS 2011.

<sup>11</sup>The survey questions that teachers use to evaluate literacy skills differ across studies. Figures B.3 and B.4 show the corresponding gradients by database and grade level. Refer to the figure notes for details on the specific survey questions.

<sup>12</sup>ECLS-K 1999 and 2011 include an index for children's "approaches to learning". Teachers report how often children keep belongings organized, show eagerness to learn new things, work independently, easily adapt to changes in routine, persist in completing tasks, pay attention well, and follow classroom rules.

<sup>13</sup>Following standard IRT procedures, the assessment is carried out in two stages. First, children answer routing questions that span a range of difficulty levels. Second, children are selected to take a low-, middle-, or high-difficulty test depending on their demonstrated ability level in the routing stage.

<sup>14</sup>Most countries test children in a single grade level—typically fourth grade—but in the 2006 wave, Iceland and Norway tested children in both fourth and fifth grade.

any of the other four plausible values instead (see Figure A.6 in the main appendix).

### B.3.4 Homework help

Our primary variable of interest for parental time investment is the extent to which someone in the home helps children with homework or schoolwork.

**ECLS-K** Information about homework help was obtained through parent interviews conducted toward the end of each school year, typically between April and June. In ECLS-K:1999, the wording of the question varied across waves with respect to who was helping the child and what subject they were helping with. By contrast, the question did not have a subject-specific focus in ECLS-K:2011, and the phrasing remained unchanged across waves. Specifically, parents were asked:

*During this school year, how often...*

*...did you help {child} with his/her homework? (ECLS-K:1999, grade 1)*

*...have you or any of the people we just mentioned helped {child} with his/her reading, language arts, or spelling homework? (ECLS-K:1999, grade 3)*

*...did someone help {child} with his/her reading, language arts, or spelling homework? (ECLS-K:1999, grade 5)*

*...did you or someone else help {child} with his/her homework? (ECLS-K:2011, grades 1, 2, 4, and 5)*

The possible responses were (1) never, (2) less than once a week, (3) one to two times a week, (4) three to four times a week, or (5) five or more times a week. We standardize the responses by database and wave.

**PIRLS** Information about homework help was obtained through the Home Questionnaire. The question changed slightly between 2006 and 2011, shifting focus from reading help to general help. Specifically, parents were asked:

*How often do you or someone else in your home do the following things with your child?*

*...Help my child with reading for school? (PIRLS 2006)*

*...Help my child with his/her schoolwork? (PIRLS 2011)*

The possible responses were (1) every day or almost every day, (2) once or twice a week, (3) once or twice a month, or (4) never or almost never. We reorder the responses from least to most compensatory, then standardize by database, country, and – when relevant – grade level. When standardizing, we apply the PIRLS sampling weights to ensure that the computed statistics are nationally representative.

### B.3.5 Skills practice

Our secondary variable of interest for parental time investment is the extent to which parents practice basic skills with their children.

**ECLS-K** Information about the frequency of skills practice was obtained through a parent interview. The wording of the question remained the same between studies and grade levels:

*In a typical week, how often do you or any other family member do the following things with {child}:*

*...Practice reading, writing or working with numbers?*

If parents asked for clarification, the interviewer explained that they should include time spent on homework, reading a calendar, or practicing in an exercise book or workbook. The possible responses were (1) not at all, (2) once or twice a week, (3) three to six times a week, or (4) every day. We standardize the responses by database and wave.

**PIRLS** The Home Questionnaire included a question on the frequency of skills practice in 2011, but not 2006. The wording of the question was:

*How often do you or someone else in your home do the following things with your child:*

*...Help my child practice his/her reading?*

*...Help my child practice his/her math skills?*

The possible responses were (1) every day or almost every day, (2) once or twice a week, (3) once or twice a month, or (4) never or almost never. Parents answered separately for reading and math skills. We re-order the responses from least to most compensatory and take the average over the two subjects for consistency with ECLS-K, which combined practice with reading and numbers in the same question. We standardize the average by country, applying PIRLS sampling weights to ensure that the computed statistics are nationally representative.

### B.3.6 Class size

Our primary variable of interest for school investment is the size of the class in which children are enrolled.

**ECLS-K** Teachers were asked about the number of students in their class who belong to a certain sex, ethnicity, and, in some waves, age level. Class size can therefore be calculated by summing up the reported numbers by sex, ethnicity, or age. The resulting sum might differ across characteristics. Such inconsistencies are, however, relatively infrequent, occurring for about 5% of the observations.

To deal with inconsistent responses, we derive a composite variable for class size following guidelines from the ECLS-K User's Manual (see Section 7.4.3.2 in [National Center for Education](#)

Statistics, 2004). When at least two of the sums by sex, ethnicity, and age level matched, we set the composite variable equal to the matching value. In cases where there were no matching sums, we set the composite variable equal to the sum over both sexes. If that data was missing, we used the sum over all age levels. If that data was also missing, we used the sum over all ethnicities.

In third to fifth grade, inconsistencies may also arise because reading, math, and science teachers completed separate questionnaires.<sup>15</sup> We use the data reported by the reading teacher whenever possible (approximately 99% of cases) and supplement it with data from the math or science teacher if data from the reading teacher is missing. Our focus on the reading teacher is motivated by the fact that our international data come from a survey of reading teachers. Additionally, we have data from the reading teacher for the full sample of children, whereas we have data from either the math or science teacher for random subsamples.

**PIRLS** In both 2006 and 2011, teachers were asked to fill in a number in response to the following question: “*How many students are in this class?*” After data collection, the variable was top-coded at 60 students. In a small number of cases, classes are linked to multiple teachers, and the enrollment count reported by the teachers differed. We take the average of the teachers’ responses for these cases (approximately 0.28% of observations).

### B.3.7 Remedial tutoring

Our secondary variable of interest for school investment is whether children receive remedial tutoring at school.

**ECLS-K** We obtain information on tutoring from the child-level teacher questionnaire. In ECLS:K 1999, the wording of the question varied slightly across waves, but it was unchanged across waves in ECLS-K:2011. The possible formulations were as follows:

*Does this child receive instruction [and/or related services] in any of the following types of programs in your school [during the school day]:*

*...Remedial tutoring program in reading?* (ECLS-K:1999, all grades, excluding the bracketed text in kindergarten and first grade)

*Does this child receive, or has he/she received during the school year, instruction in any of the following types of programs in your school:*

*...Remedial tutoring or remedial program in reading/language arts?* (ECLS-K:2011, all grades)

For our analysis, we create an indicator equal to one if the teacher responds “yes” to the question.

---

<sup>15</sup>These may not be true inconsistencies if students are sorted into different classes for each subject. It is, however, common for the same group of students to take all of their classes together, even if they have multiple teachers.

**PIRLS** We cannot study this outcome using PIRLS data because teachers do not respond to child-level questionnaires.

### **B.3.8 Educational expectations**

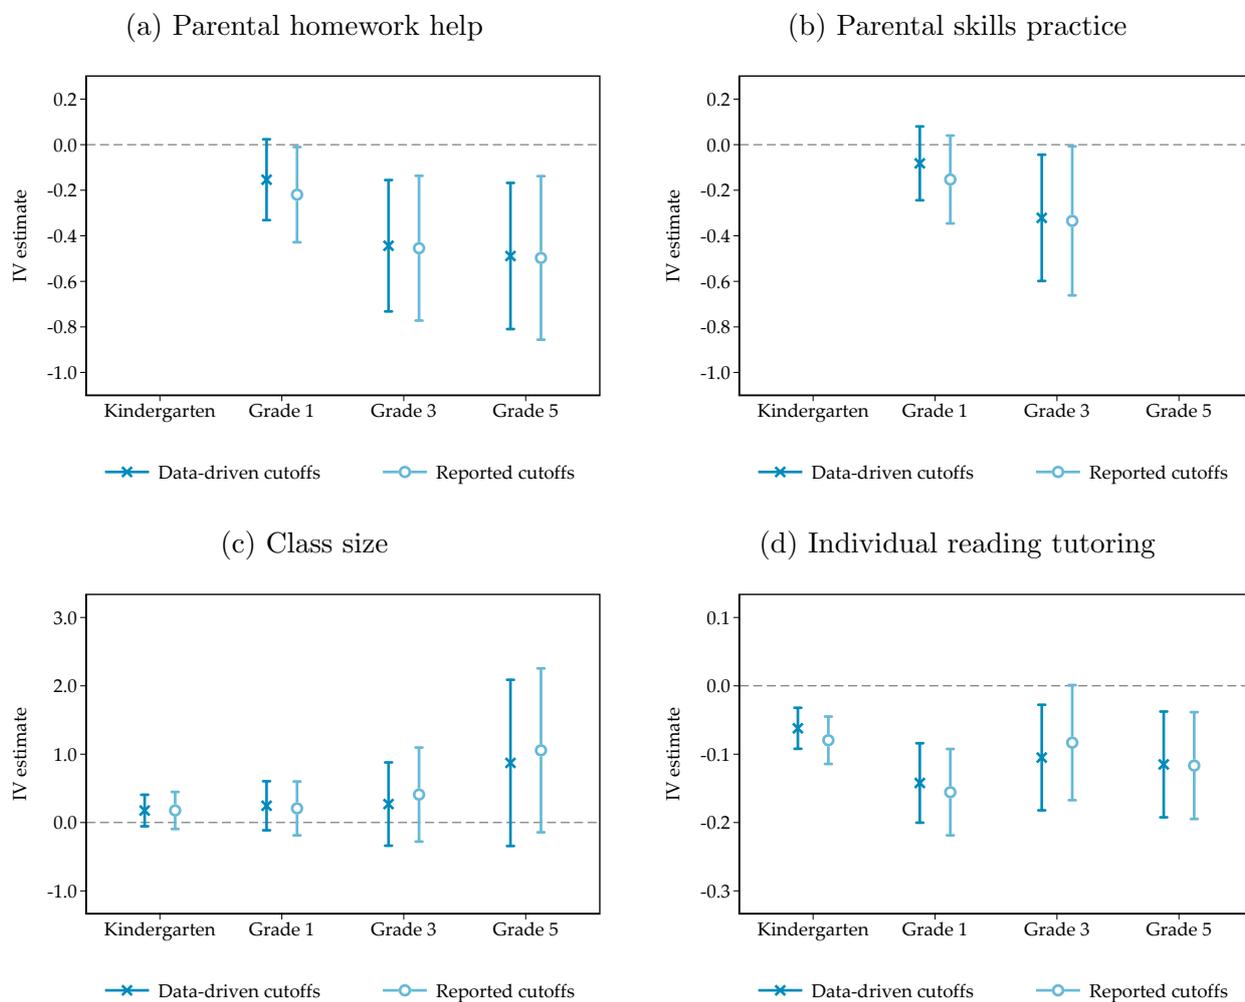
To assess whether parents and schools update their priors as information on student performance is revealed, we study teacher and parental expectations regarding children’s future educational attainment. We are primarily interested in the dynamic effects over grade levels, and hence, our main analysis focuses on the ECLS-K panel data for this analysis. Nevertheless, for completeness, we also present results based on cross-sectional data from PIRLS 2011 in the main appendix.

**ECLS-K** The parent interview included the following question on educational expectations: *“How far in school do you expect (child) to go?”*. The question was asked in kindergarten, first grade (ECLS-K:1999 only), third grade, and fifth grade. For our analysis, we create an indicator equal to one if parents responded that they expect their child to finish a college degree or higher (i.e., at least four years of tertiary education). There is no data on teacher expectations in ECLS-K:1999. However, in the 2011 study, the teacher questionnaire included the following question for first and second graders: *“How far in school do you think this child will go”?* For consistency with our parental expectations variable, we create an indicator equal to one if the teacher responds that they expect the child to complete a four- or five-year college degree or higher.

**PIRLS** The Home Questionnaire did not include a question on parental expectations in the 2006 wave. In 2011, parents were asked: *“How far in his/her education do you expect your child to go?”* We create an indicator variable equal to one if parents responded that they expect their child to finish ISCED Level 5A or beyond (i.e., tertiary education). There is no child-level data on teacher expectations in either 2006 or 2011 because the teacher questionnaire collects information on class characteristics rather than individual students.

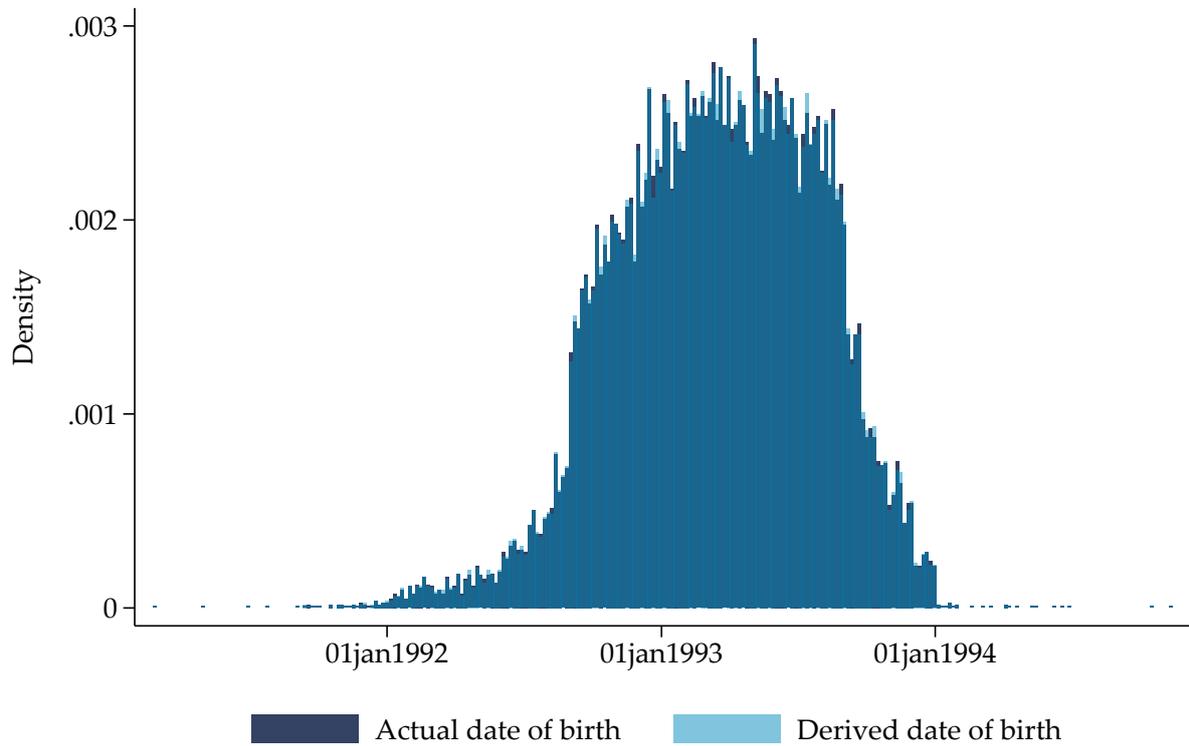
## B.4 Supplementary figures and tables

Figure B.1: Robustness check – comparison of IV estimates for ECLS-K:1999 using data-driven and reported SSA cutoff dates



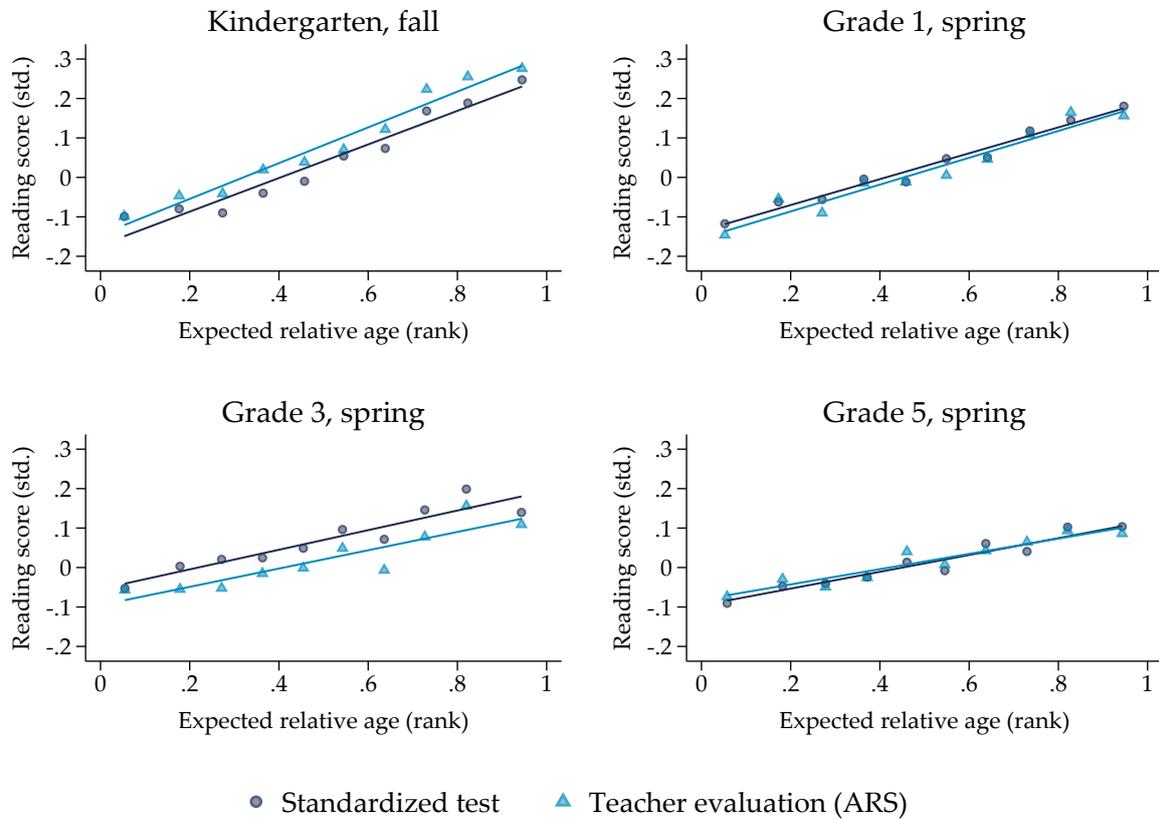
*Notes:* This figure uses data from ECLS-K:1999 to illustrate that our main IV estimates are unchanged when we calculate our instrument (children’s expected relative age rank) using the SSA rules reported by school administrators instead of the SSA rules derived using our data-driven approach. The bars show 95% confidence intervals when standard errors are clustered by expected date of birth.

Figure B.2: Distribution of actual and derived date of birth in ECLS-K:1999



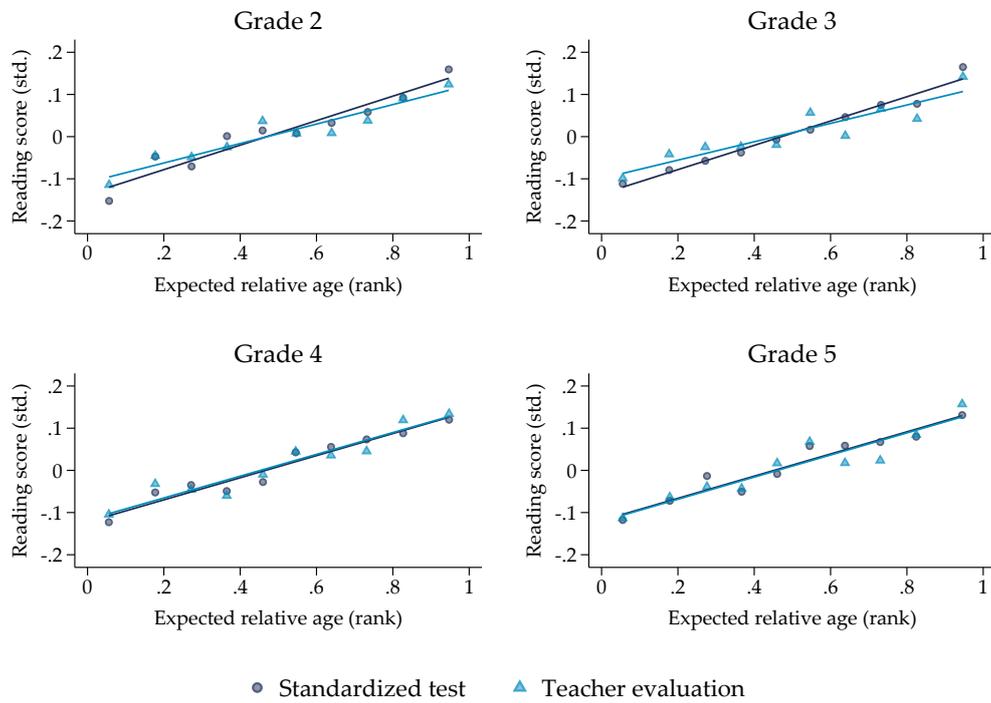
*Notes:* This figure uses data from ECLS-K:1999 to illustrate that the method we use to derive children's date of birth in the 2011 study produces a date-of-birth distribution (lighter shading) that is nearly identical to the distribution of children's actual date of birth (darker shading).

Figure B.3: Relationship between expected relative age rank and different forms of reading assessments in ECLS-K:1999



*Notes:* Data comes from ECLS-K:1999. Teachers were asked to evaluate multiple aspects of children’s literacy skills – for instance, their ability to name letters of the alphabet or read simple books independently relative to similarly aged children. NCES used the teacher responses to compute an Academic Rating Score (ARS). We standardize the reported ARS score to have mean 0 and standard deviation 1.

Figure B.4: Relationship between expected relative age rank and different forms of reading assessments in ECLS-K:2011



*Notes:* Data comes from ECLS-K:2011. Teachers are asked to rate children's reading skills based on the curriculum standards for his/her grade level. The possible responses are: below grade level, about on grade level, and above grade level. We standardize the responses to have mean 0 and standard deviation 1.

Table B.1: School-starting-age (SSA) rules in the countries in our sample

Country	SSA rule	PIRLS Participant	
		2006	2011
Australia	Varies by region	No	Yes <sup>†</sup>
Austria	Age 6 by September	Yes	Yes
Belgium	Age 6 by January	Yes	Yes
Bulgaria	Age 7 by January	Yes	Yes
Canada	Varies by region	Yes <sup>†</sup>	Yes <sup>†</sup>
Croatia	Age 6 by April	No	Yes
Czech Republic	Age 6 by September	No	Yes
Denmark	Age 7 by January	Yes	Yes
England	Age 5 by September	Yes	Yes <sup>‡</sup>
Finland	Age 7 by January	No	Yes
France	Age 6 by January	Yes	Yes
Germany	Varies by region	Yes	Yes <sup>‡</sup>
Hong Kong	Age 6 by January	Yes	Yes
Iceland	Age 6 by January	Yes	No
Iran	Age 6 by Mehr <sup>*</sup>	Yes <sup>‡</sup>	Yes
Israel	Age 6 by Tevet <sup>*</sup>	Yes <sup>‡</sup>	Yes
Italy	Age 6 by January	Yes	Yes
Latvia	Age 7 by January	Yes	No
Lithuania	Age 7 by January	Yes	Yes
Luxembourg	Age 6 by September	Yes	No
Malta	Age 5 by January	No	Yes
Netherlands	Age 6 by October	Yes	Yes
Northern Ireland	Age 4 by July	No	Yes
Norway	Age 6 by January	Yes	Yes
Poland	Age 7 by January	Yes	Yes
Portugal	Age 6 by January	No	Yes
Scotland	Age 5 by March	Yes	No
Singapore	Age 7 by January	No	Yes
Slovak Republic	Age 6 by September	Yes	Yes
Slovenia	Age 6 by January	Yes	Yes
Spain	Age 6 by January	Yes	Yes
Sweden	Age 7 by January	Yes	Yes
Taiwan	Age 6 by September	Yes	Yes
United States	Varies by region	Yes <sup>‡</sup>	Yes <sup>‡</sup>

*Notes:* The rules pertain to students' age of entry to primary school; e.g., "Age 6 by September" means that a child must be six years or older on August 31 to start school. We derive the information from questionnaires completed by the PIRLS National Research Coordinator in each country. For details, see question ACQ02 of the Curriculum Questionnaire in the [PIRLS 2006 User Guide](#) and Appendix C.1 in the [PIRLS 2011 User Guide](#).

\* Tevet begins sometime in December. Mehr typically begins on September 22 or 23.

<sup>†</sup> PIRLS data from certain regions is excluded from our analysis. See Table B.2.

<sup>‡</sup> PIRLS data from the entire country is excluded from our analysis. See Table B.2.

Table B.2: Sample selection – Regions dropped from the PIRLS data

<b>Reason for exclusion</b>	<b>Countries/regions</b>
Expected relative age cannot be coded properly due to mid-month birthdate cutoff <sup>†</sup>	Iran (2006); Israel (2006)
Reported SSA rule is ambiguous, not uniform across local education authorities, and/or not verifiable in the observed date-of-birth distribution (i.e., first-stage relationship is weak or non-existent)	Australia—New South Wales; Azerbaijan; Botswana; Canada—Alberta; Colombia; Georgia; Honduras; Hungary; Indonesia; Ireland; Kuwait; Macedonia; Moldova; Morocco; New Zealand; Oman; Qatar; Romania; Russia; Saudia Arabia; South Africa; Trinidad and Tobago; United Arab Emirates
Home Questionnaire was not distributed to parents	England (2011); United States

*Notes:* <sup>†</sup>In PIRLS 2006, we do not observe children’s exact date of birth – only the month and year. Thus, in countries where the date of birth stipulated by the SSA rule falls in the middle of the month, we cannot code children’s expected date of birth correctly.

Table B.3: School starting age rules in Australia for the 2011 cohort

State/Territory	SSA rule	PIRLS Participant	
		2006	2011
Australian Capital Territory	Age 6 by May	No	Yes
New South Wales	Age 6 by August	No	Yes
Northern Territory	Age 6 by July	No	Yes
Queensland	Age 6 by January*	No	Yes
South Australia	Age 6 by May	No	Yes
Tasmania	Age 6 by January	No	Yes
Victoria	Age 6 by May	No	Yes
Western Australia	Age 6 by July	No	Yes

*Notes:* The reported rules correspond to the age of entry to primary school (Year 1). Information on cutoff dates is derived from government documents provided by the [Queensland Department of Education](#) (Disclosure Log 340/5/2044). The symbol \* indicates that the cutoff date changed after our study period. In the data, we identify states and territories through the explicit stratification variable, IDSTRATE. The codes are listed in Supplement 4, Exhibit S4.1 of the [PIRLS 2011 User Guide](#).

Table B.4: School starting age rules in participating Canadian provinces

Province/Territory	SSA rule	PIRLS Participant	
		2006	2011
Alberta	Varies locally	Yes <sup>†</sup>	Yes <sup>†</sup>
British Columbia	Age 6 by January	Yes <sup>†</sup>	Yes
Manitoba	Age 6 by January	No	Yes
New Brunswick	Age 6 by January	No	Yes
Newfoundland and Labrador	Age 6 by January	No	Yes
Nova Scotia	Age 6 by October	Yes <sup>†</sup>	Yes
Ontario	Age 6 by January	Yes <sup>†</sup>	Yes <sup>†</sup>
Quebec	Age 6 by October	Yes <sup>†</sup>	Yes <sup>†</sup>
Saskatchewan	Age 6 by January	No	Yes

*Notes:* Information is derived from the PIRLS 2006 and 2011 Curriculum Questionnaires. The symbol † indicates that the region was a benchmarking participant. Data for benchmarking participants is recorded in province-level databases. Abbreviations are used in the file names to indicate which province the data comes from: CAB (Alberta); CBC (British Columbia), CNS (Nova Scotia); COT (Ontario); and CQU (Quebec).

Table B.5: Abbreviated list of cutoffs for ECLS-K data-driven approach

<b>Panel (a): ECLS-K:1999</b>		
<b>Combined cutoff</b>	<b>Other dates included</b>	<b>Nr. of states</b>
June 1, 1993	—	1
August 1, 1993	—	1
August 15, 1993	—	1
September 1, 1993	August 31 & September 2, 1993	22
September 15, 1993	September 10, 1993	3
October 1, 1993	September 30, 1993	6
October 15, 1993	—	3
December 1, 1993	December 2, 1993	3
January 1, 1994	December 31, 1993	5
<b>Panel (b): ECLS-K:2011</b>		
<b>Combined cutoff</b>	<b>Other dates included</b>	<b>Nr. of states</b>
August 1, 2005	—	2
August 15, 2005	—	3
September 1, 2005	August 31 & September 2, 2005	23
September 15, 2005	September 10, 2005	3
October 1, 2005	September 30, 2005	6
October 15, 2005	—	2
December 1, 2005	December 2, 2005	3
January 1, 2006	December 31, 2005	4

*Notes:* The first column lists the cutoff dates that we iterate over when implementing our data-driven approach to identify the SSA rule for each school. We combine cutoffs that are too close in time to distinguish using our algorithm: for example, cutoffs on Aug. 31, Sep. 1, and Sep. 2 are combined into a Sep. 1 cutoff. The third column lists the number of states in which the rule applies. The number of states does not sum to 50 because some states do not have a uniform SSA rule. We obtained information on the legislated state cutoffs from [Elder and Lubotsky \(2009\)](#) and [Education Commission of the States \(2010\)](#).

# References

- Education Commission of the States**, “ECS State Notes: State Kindergarten Statutes – 50-State Reports,” December 2010.
- Elder, Todd E. and Darren H. Lubotsky**, “Kindergarten Entrance Age and Children’s Achievement: Impacts of State Policies, Family Background, and Peers,” *The Journal of Human Resources*, 2009, 44 (3), 641–683.
- Foy, Pierre and Kathleen T. Drucker, eds**, *PIRLS 2011 User Guide for the International Database*, TIMSS & PIRLS International Study Center, Lynch School of Education, Boston College, 2013.
- Labrecque, Mélanie, Maria Chuy, Pierre Brochu, and Koffi Houme, eds**, *PIRLS 2011 Canada in Context*, Council of Ministers of Education, Canada, 2012.
- National Center for Education Statistics**, *User’s Manual for the ECLS-K Third Grade, Public-Use Data File and Electronic Code Book, Early Childhood Longitudinal Study, Kindergarten Class of 1998-99, NCES 2004-001*, Institute of Education Sciences, U.S. Department of Education, 2004.
- Schwandt, Hannes and Amelie Wuppermann**, “The Youngest Get the Pill: ADHD Misdiagnosis in Germany, Its Regional Correlates and International Comparison,” *Labour Economics*, 2016, 43, 72–86.