Distorted Beliefs and Parental Investment in Children

Josh Kinsler, University of Rochester

Ronni Pavan, Royal Holloway, University of London

Richard DiSalvo, University of Rochester

December 8, 2014

Abstract

Parental investments in early childhood have been shown to have a large impact on skill acquisition. In this paper, we examine how beliefs about a child's relative ability influences investment and how these beliefs are determined. Using data from the ECLS-K, we first show that a parent's beliefs about her child's ability relative to children of the same age is heavily influenced by her child's ability relative to children in the same school even after controlling for a measure of overall ability. We then show that this local distortion in beliefs affects remedial parental investments such as helping with homework or hiring a tutor. Building off our descriptive findings, we then develop and estimate a structural model of parental investment that incorporates uncertainty, learning, and local distortions. We estimate the model using indirect inference and then perform a series of counterfactuals where we eliminate the local information distortions or, alternatively, eliminate the ability stratification implied by the observed sorting into schools. In either case we find that investment and achievement rise by a considerable amount for students at the bottom of the ability distribution. The reason for this is that with more accurate information parents of children in relatively low achieving schools realize how far behind their children actually are.

1 Introduction

There is a recent and growing literature demonstrating that parental investment is an important input in the production of adolescent skill.[1] This research has sparked renewed interest in understanding the determinants of parental investment during childhood. Models of parental investment typically focus either on the impact of credit constraints or the tradeoff between goods and time investments.[2] Relatively little attention has been paid to the role of uncertainty, information, and learning in parental decision making. This is in stark contrast to much of the recent literature on own human capital investments, where imperfect information about ability or returns plays a central role.[3]

In this paper, we examine the determinants of parental beliefs about child ability and how these beliefs influence subsequent investment. To do this we utilize data collected in the Early Childhood Longitudinal Study, Kindergarten Class of 1999 (ECLS-K). Surveyed parents are asked to compare their child's ability to the ability of similarly aged children. A key finding of our paper is that a parent's beliefs about her child's ability relative to children of the same age is heavily influenced by her child's ability relative to children *in the same school*. In other words, parents of a student who is above average in their school are more likely to believe that their child is above average overall. To the extent that parents care about their child's rank in the overall ability distribution, say for attending college, the tendency to rely on local comparisons can have a significant distortionary impact on investment behavior. Moreover, if parental investment is compensatory then the bias in parental beliefs has the potential to exacerbate gaps in student outcomes. Parents of children attending schools with relatively low average ability will invest less than the

[1]For example, Cunha *et al.* (2010) find that measured parental investment accounts for 15% of the variation in educational attainment. Carneiro & Heckman (2003), Cunha & Heckman (2008), and Todd & Wolpin (2007) provide additional supporting evidence on the importance of early childhood investment. Heckman & Mosso (2014) provide a nice summary of the literature.

[2]Caucutt & Lochner (2012) and Cunha (2013) estimate dynamic models of investment focusing on the role of credit constraints. These papers build on earlier work by Becker & Tomes (1986) who were interested in a similar question. Boca *et al.* (2014) and Bernal (2008) estimate dynamic models of investment focusing on the labor supply and time allocation problem facing parents.

[3]The role of learning and uncertainty has been explored in the following human capital contexts: college dropout (Stinebrickner & Stinebrickner (2009), Arcidiacono *et al.* (2014)), major choice (Kinsler & Pavan (2014), Wiswall & Zafar (2014)), and occupational choice (Antonovics & Golan (2012), Sanders (2014)). This list is not meant to be exhaustive, but simply to illustrate the prevalence of the subject.

optimal level since they perceive that their children are better off than they actually are. The opposite pattern would exist in relatively high achieving schools.

Two recent papers also explore the idea that parents might be misinformed and thus choose investments sub-optimally. Cunha *et al.* (2013) surveys a sample of socioeconomically disadvantaged, pregnant African American women and elicits their subjective expectation about the elasticity of child development with respect to investment. The median reported elasticity is between 4% and 19%, while estimates from the CNLSY/79 indicate an elasticity between 21% and 36%. If the median mother in the survey were given the objective elasticities, investment is estimated to increase between 4% and 24% with a subsequent increase in cognitive skills between 1% and 5%. Using data from a field experiment in Malawi, Dizon-Ross (2013) finds that parents' perceptions of their children's recent achievement diverges substantially from children's true recent achievement, with an average gap between the two a full standard deviation. Providing parents with accurate information causes them to re-allocate their educational investments.

While we are also interested in understanding how investment behavior might change if parents were better informed, our approach is somewhat different. We aim to understand the nature of parental bias and the possible sources of the bias. This is critical for designing and implementing policies aimed at improving parental information over multiple periods during childhood. As an example, imagine that the source of bias in parental beliefs stems from continuous interactions with family and friends. Revealing the true level of a child's skills once might be an ineffective policy for changing parental behavior in the long term since the bias would likely reproduce itself overtime as a result of a continuous exposure to inaccurate information. In our setting, parents and students repeatedly interact with the local ability distribution as they progress through school. Thus, removing the distortionary impact of local rank on parental beliefs about the overall rank of their child likely requires repeated intervention.

To illustrate the influence of a child's local relative ability on parental beliefs about a child's overall relative ability, we first estimate a series of reduced-form regressions describing parental beliefs. The key regressors are a student's standardized test score and the difference between a student's test score and the average test score in a student's school. We find that the deviation of a child's test score from the school average has a large and statistically significant impact. This suggests that parents' beliefs about their child's ability relative to similarly aged children are systematically biased towards the child's ability relative to children in the same school. We provide additional evidence to help rule out alternative explanations.

Having established that parental beliefs about a child's overall relative ability are influenced by a child's local relative ability, we next show that parental investment is strongly related to beliefs about overall relative ability. Our focus is on remedial investment, which includes helping with homework and tutoring, since we think these investments are most directly related to the academic achievement of children. We find that parents who believe their child is above average invest 10% to 20% of a standard deviation less than all other parents. Moreover, remedial investment is more responsive to parental beliefs than other types of investments, such as reading, playing, or singing songs.

One question the belief and investment regressions do not address is how parents form expectations. In other words, what signals are parents receiving such that the local relative ability plays such a prominent role in beliefs about overall relative ability? We investigate teachers as one possible conduit. We show that teacher beliefs about a child's ability relative to other children in the same grade are heavily influenced by the difference between a child's test score and the school average test score. Moreover, the evidence suggests that these teacher ratings influence parents' beliefs about their child's rank.

Our descriptive results indicate that parents have distorted beliefs about the ability of their children relative to other children of the same age, these beliefs significantly influence investment, and that teachers are one potential channel through which parent's receive biased information. Using these descriptive results as a guide, we then develop and estimate a structural model of parental beliefs, information, and investment. The benefit of the structural model is that it allows us to perform the type of counterfactual analyses discussed earlier. In particular, we are able to examine how parental investment and student outcomes change when we correct the distortions in parental beliefs. We are also able to compare the impact of one-time interventions with repeated interventions.

The basic components of the model are as follows. Each child is characterized by an unobserved relative ability in kindergarten, 1st, and 3rd grade. At the end of each grade, parents receive signals about their child's overall relative ability, some of which will be biased towards the child's local relative ability. Parents use the signals to update their beliefs about their child's overall relative ability, unaware of any bias, and then choose investment in order to maximize the probability that their child will obtain a four year college degree. Following investment, a child's relative ability evolves according to a log-CES production function. We estimate the model using indirect inference, targeting moments based on regressions similar to those described earlier. We experiment with various weighting matrices, including the identity and optimal weighting matrix.

Estimates from the model indicate that correcting the distortions in parental beliefs would lead to large changes in behavior. Using the optimal weighting matrix, we predict that parents of students in the bottom 10% of the ability distribution would increase investment in 1st and 3rd grade by 20% and 50% of a standard deviation when they receive unbiased signals about their child's ability relative to other children of the same age. Test scores for this group subsequently increase by approximately 10% of a standard deviation in 3rd grade. Predicted changes in parental behavior and student outcomes at the bottom of the ability distribution are even larger when we use the identity matrix for estimation. Regardless of how we weight the underlying moments in estimation, parents of students at the top of the ability distribution reduce investment slightly with no appreciable impact on 3rd grade test scores.

The remainder of the paper is as follows. In Section 2 we discuss the ECLS-K data in detail. In Section 3 we present reduced-form evidence on distorted beliefs and parental investment. We develop a formal model of parental beliefs, information, and investment in Section 4. In Section 5 we discuss our estimation approach and provide parameter estimates. We perform counterfactual analyses in Section 6 and conclude in Section 7.

2 Data

We use the Early Childhood Longitudinal Study, Kindergarten Class of 1999 (ECLS-K) to study parental beliefs, investment, and student outcomes. The ECLS-K is a longitudinal study that surveys a nationally representative sample of parents, children, teachers, and school administrators in the spring of kindergarten, 1st, 3rd, 5th, and 8th grades.[4] 21,409 children distributed across 1,018 schools are included in the initial kindergarten sample. Information about a child's home, school, and classroom environments is collected. We focus our analysis on data collected prior to 5th grade. In the following paragraphs we discuss the key variables used in our analysis and describe the process by which we arrive at our estimation sample.

In each round of the survey, students are assessed in reading and math. The ECLS-K assessed skills that are typically taught and developmentally important, and the assessment frameworks were derived from national and state standards. The cognitive assessments were two-stage adaptive tests; all children began a subject area test with a routing test, which was then followed by a second-stage form. The two-stage, adaptive assessment format helped ensure that children were tested with a set of items most appropriate for their level of achievement and minimized the potential for floor and ceiling effects. We standardize the Item Response Theory Scale Scores from the reading and math assessments and utilize these as unbiased measures of a child's relative ability.[5]

Parental beliefs about their child's ability relative to other children of the same age are elicited in the fall of kindergarten and in the spring of 1st and 3rd grade. The precise wording of the question is as follows: "Does your child learn, think, and solve problems better, as well, slightly less well, or much less well than other children his/her age?" In the fall of kindergarten, 92% of parents respond that their child performs better or as well

[4]There is an additional survey in the fall of kindergarten and for a subsample of the original data a survey in the fall of 1st grade.

[5]Item Response Theory uses the pattern of right, wrong, and omitted responses to the items actually administered in a test and the difficulty, discriminating ability, and "guess-ability" of each item to place each child on a continuous ability scale. The items in the routing test, plus a core set of items shared among the different second stage forms, made it possible to establish a common scale. It is then possible to estimate the score the child would have achieved if all of the items in all of the test forms had been administered.

as other children of the same age. In the spring of 1st and 3rd grade, parents are also asked to compare the math and reading skills of their child to the math and reading skills of the other children in their child's class. Here parents are asked, "Compared to other children in your child's class, how well do you think he/she is doing in school this spring in math? Do you think he/she is doing much worse, a little worse, about the same, a little better, or much better?" A similar question is asked for reading. These classroom based questions are useful for demonstrating that parents understand the difference between local and global comparisons and utilize different reference points to assess their child's skills.

Kindergarten, 1st, and 3rd grade teachers are also asked to assess the math and reading skills of the surveyed children. Similar to the parent questions, teachers are asked to compare the child's math and reading academic skills to other children of the same grade level. The choices available to the teacher are: far below average, below average, average, above average, and far above average.[6]

Along with beliefs, parental investment is the other essential variable for our model. We calculate "remedial" investment in 1st and 3rd grade as the primary factor of three underlying variables. The first variable is the number of times per week a parent helps their child with homework during the past school year. This variable captures the degree to which parents engage and assist their child in learning school material. The second variable is the ratio of the number of times per week the parent helps the child with homework to the number of times a child does homework at home.[7] Here we attempt to account for the fact that there may be variability in the amount of homework that children receive. The final variable is whether the child is tutored on a regular basis by someone other than a family member. The largest factor loadings are associated with the questions related to homework assistance.

[6]Teachers are also asked to rate specific math and reading skills of the the child on a five point scale. Examples of specific math skills in first grade are whether a child understands place value and uses a variety of strategies to solve math problems. The rating scale reflects the degree to which a child has acquired and/or chooses to demonstrate the targeted skills, knowledge, and behaviors. The ECLS-K combines these specific ratings using IRT to generate a single rating. These math and reading ratings correlate with the math and reading ratings based on comparisons to children in the same grade at a level of 0.72 and 0.82 respectively.

[7]In 3rd grade we use this ratio for both math and reading in our factor model.

The ECLS-K has additional variables that can be interpreted as parental investments. In the fall of kindergarten parents are asked how many times per week they tell stories, sing songs, read books, play games, play sports, do arts and crafts, and do science projects with their children. Starting in first grade, parents are also asked whether the child takes music, art, or drama and how often the family eats together, goes to museums, and receives newspapers or magazines. We combine these measures of investment into a separate factor which we denote "activities" investments. In Section 3.3 we examine the sensitivity of both "remedial" and "activity" type investments to parental beliefs.

One final variable that is important for our structural model is parental beliefs about how far they expect their child to go in school. This variable is strongly related to parents' beliefs about the ability of their child relative to children of the same age. In the model we assume parents choose investment to maximize the likelihood that their child obtains a four-year college degree and we tie this probability to the data using parent reports. One concern with this might be that parents of children at the very bottom of the ability distribution will have no incentive to invest since the likelihood of graduating from a four-year college is very small. However, Figure 1 indicates that this is not the case. First, the fraction of parents that believe their child will complete a four year college degree is strictly increasing in the decile of math score in third grade, as we would expect. Second, a large fraction (more than 60%) of the parents of even the lowest decile of the students distribution believes that their child will graduate from college.

In addition to beliefs and investments, the survey also contains basic demographic and socioeconomic variables, such as race, gender, and family income in the fall of kindergarten. Another important feature of the data is the ability to group respondents together in schools. In the fall of kindergarten, we observe approximately twenty-one survey respondents per sampled school. This allows us to create proxies for the average math and reading ability in each school.

While the survey data is incredibly rich, the challenge in working with the ECLS-K is the high level of attrition. This is particularly problematic in our setting since we want to maintain a reasonable number of students in each school so that our proxies for school ability are still informative. In the fall of kindergarten, there are 21,409 sampled children. We start by dropping 149 students who lack a valid school identifier in kindergarten and an additional 1,650 students who attend kindergarten in a school with fewer than 5 students. Then, between the fall of kindergarten and the spring of 1st grade approximately 3,400 students attrit, and an additional 2,400 attrit between the spring of 1st and 3rd grade.

In addition to students completely disappearing from the sample, we also drop students who splinter off from their original school. As an example, consider a school that has 25 surveyed children in the fall of kindergarten. Now imagine 18 of these students matriculate to the same 1st grade school, 5 students splinter off to a second 1st grade school, and 2 students splinter off to a third 1st grade school. Because we want to focus on relatively large schools we drop the 7 students that splinter off from the main group. We do this for both the transition to 1st grade and the transition to 3rd grade. In addition, we drop any student who attends a school where the largest matriculating group is less than 50% of the total. Between kindergarten and first grade and 1st and 3rd grade we lose about 2,200 and 1,600 additional students as a result of this splintering behavior.

Our final sample contains 19,576, 14,940, and 11,253 students in the fall of kindergarten, spring of 1st grade, and spring of 3rd grade respectively. Table 1 provides means for the key variables discussed above for each grade in our sample. The first few rows of the table indicate that attrition is not entirely random since the sample becomes increasingly white and wealthier as measured by income in the fall of kindergarten. Also, the number of students per school declines considerably as a result of attrition. The parental belief variables indicate a significant skewness. More than 30% of the sample think their child thinks and solves problems better than other children his/her age, while only 7% think their child thinks and solves problems slightly less well or much less well than other children. Similar patterns are observed when parents are asked to compare their child to other children in the same class. Consistent with these skewed beliefs, we see that approximately 80% of parents think their child will obtain a 4-year college degree. The final few rows show the variability across households in helping with homework. Approximately 25% of parents help their child every day, while 5% never help their child with homework.

3 Evidence on Distorted Beliefs and Investment

In this section we present reduced-form evidence supporting our hypotheses regarding parental beliefs and investment. We first illustrate that parental beliefs about a child's ability relative to similarly aged children is distorted by a child's ability relative to his/her class. Next, we show that teacher beliefs about a child's ability is similarly distorted and that teacher beliefs significantly influence parental beliefs. This suggests an important channel through which parents may receive biased information. We then demonstrate that parental investment is strongly related to beliefs about a child's ability relative to similarly aged children. Finally, we present evidence on the productive effects of investment.

3.1 Distorted Beliefs

In each round of the survey parents are asked to compare their child's ability to learn, think, and solve problems to children of a similar age. The answer to this question is our main outcome variable. Parents are given four options, however, almost all parents respond that their child is either better or as good as similarly aged children. Thus, for most of this section we treat parental beliefs as if they are binary, with a one indicating that they believe their child is above average.

Table 2 illustrates the relationship between parents' beliefs about whether their child is above average and math test scores using a linear probability model. We focus on students who are in 1st and 3rd grade since this allows us to make direct comparisons to parental beliefs about their child's ability relative to children in the same class. The first column indicates that a one-standard deviation increase in the ECLS-K math assessment increases the probability that parents think their child is above average by 16 percentage points. In the second column we incorporate an additional regressor, the difference between a student's math score and the average math score in a student's school. The impact of the math score alone declines so that a one-standard deviation increase implies an 11 percentage point increase in the probability that parents respond that their child is above average. A one-standard deviation increase in the difference between the child's score and the school average increases the likelihood of an above average report by 6 percentage points. Thus, local relative ability is exuding influence on parental beliefs about overall relative ability.

One concern regarding the above result is that parents may simply be misinterpreting the question. If parents believe they are being asked to compare their child to other kids in the child's class, we would expect the deviation from the school average to matter. However, we can investigate this directly since parents are also asked about their child's math and reading skills compared to other children in their class. The third and fourth columns of Table 2 change the dependent variable to an indicator for whether parents think their child is much better than the other children in his/her class in math and reading respectively. In these two regressions, it is only the difference between the test score and the school average score that has a meaningful impact on beliefs. This suggests that parents understand that they are being asked two different questions and utilize different reference points to assess their child's skills.

When parents are asked to compare their child to children of the same age, they do so according to how well the child learns, thinks, and solves problems. This question does not map exactly to the math skills being measured by the ECLS-K and maybe this has something to do with the impact of the test score deviation from the school average. However, Table 3 illustrates that the relative importance of the test score deviation from the school average score is nearly identical when we use reading scores instead of math scores.

One potential problem affecting both the math and reading score regressions is measurement error. If a student's own test score is sufficiently noisy, then the school average might enter the belief regressions significantly since it may also be a noisy measure of a student's underlying skill. Note that if this were the case we would expect the test score deviation from the school average to negatively influence beliefs. A student who attends a school with high average test scores would tend to have higher unobserved skills. Nevertheless, we investigate the role of measurement error in our initial regressions by instrumenting for both the own score and the deviation from the school average using lags of these same variables. The idea is that the instruments will only pick up true skills as opposed to any measurement error. The first four columns of Table 4 indicate that our main findings are robust to concerns about measurement error. In fact, when we instrument the relative importance of the deviations from the school average increases.

As an additional check on the potential for measurement to be biasing our baseline estimates, we regress a direct measure of skill on test score and the deviation from the school average test score. If measurement error were driving the impact of the deviation from the school average on parental beliefs, we would expect to see a similar pattern when we use a direct measure of skill as the dependent variable. The measure of skill we utilize is whether the child reads for fun at home at least three times a week. The last two columns of Table 4 indicate that the deviation from the school average score has a negative impact on reading for fun at home. This is the opposite of what we find for parental beliefs.

The results thus far strongly suggest that parental beliefs about a child's ability relative to children of a similar age are influenced by comparisons between the child and his/her schoolmates. We interpret this as evidence that beliefs about a child's overall relative ability are biased by the local distribution of ability. However, an alternative interpretation could be that parents have heterogenous reference points for similarly aged children. For example, parents in California don't compare their child to children in Texas since they won't compete for the same colleges or future jobs. If this is the case, then a child's standing in the local ability distribution may be more important for parents than the child's placement in the overall distribution. We investigate whether this type of behavior can explain the patterns we observe by including additional test score deviations based on groupings broader than the school. The two grouping we consider are based on socioeconomic characteristics and geography. For socioeconomic groupings, we find the average test score for students of the same race and gender with similar family incomes. For groupings based on geography, we construct the average test score by census region. Table 5 illustrates that when deviations from these averages are included in our baseline regressions, the coefficients on the own score and deviation from the school average are essentially unchanged. This suggests that the school deviation is not picking up a more "local" comparison then all children of a similar age.

As a final check on the robustness of our result, we examine whether school level heterogeneity in beliefs that is correlated with average scores could be biasing our results. To do this we first estimate a model that includes an additional school characteristic, the amount of time children are expected to work on homework. The first four columns of Table 6 indicate that including assigned homework does not alter our original finding. Next, we calculate the difference between each child's test score and the average test score in his/her class. Note that this deviation is much noisier than the difference with the school average since we typically have only a handful of children in a particular class. Using this measure, however, allows us to examine the impact of the test score deviation from the class average *within* schools. The last four columns of Table 6 show that the classroom deviations are significant predictors of parental beliefs even when we control for school fixed effects. The coefficients are smaller than the coefficients on the school deviations, but again the decline is partly a result of increased measurement error.

To summarize, we find that parents are significantly more likely to report that their child thinks, learns, and solves problems better than other similarly aged children if the child's math or reading test score is higher than the average test score in their school. This result is robust to a number of plausible alternative hypotheses. We interpret the relationship between parental beliefs about overall relative ability and deviations from the school average test score as an indication that parents are biased by local information. In the next section we provide evidence that teachers may be one channel through which parents obtain biased information.

3.2 Teachers as a Source of Bias

While parents may be able to observe the skills of their child through repeated interactions, it is difficult for parents to discern the skills of other children. This is precisely why parents may have difficulty comparing their child to other similarly aged children. Parents may have a better sense for the skills of the children in the school their child attends, but it is still difficult for parents to observe this directly. However, there is a natural aggregator of information at the school level, teachers.

Teachers typically send home multiple report cards over the course of a year and may meet with parents on various occasions. Through these interactions, teachers are able to convey to parents their opinions regarding a child. If teacher beliefs about a child's ability are also in part driven by where a child falls in the school distribution, then teachers may be one source of parental biases.

As noted in the data section, teachers are also asked to compare each child to other children in the same grade. We interpret "in the same grade" similarly to the parent comparison "of the same age". Thus, we view these ratings as the teachers opinion about where the child falls in the overall ability distribution.[8] Table 7 shows how teacher ratings about a child's relative ability are influenced by test scores and test score deviations from the school average. Similar to parental beliefs, teacher ratings are also significantly influenced by a child's test score deviation from the school average. It is true for both math and reading, and is robust to controlling for lagged teacher ratings, test scores, and test score deviations from the school average.

The evidence from Table 7 indicates that teachers have similar biases to parents. Is it possible that teachers transmit their bias to parents? To investigate this idea we examine whether teacher ratings influence parental beliefs. Table 8 shows that higher teacher ratings significantly increase the likelihood that parents will say their child is above average relative to similarly aged children. This is true even when we control for parental beliefs about how the child compares to his/her classmates, lag scores and teacher ratings, and school fixed effects. The fact that parent beliefs are affected by contemporaneous teacher ratings even when we control for lags suggest that parents are using new information to update their beliefs about their child.

3.3 Beliefs and Investment

Having established that beliefs are "locally" biased, we examine whether parents act on these beliefs. In the following discussion we aim to show that parents respond to beliefs

[8]Teachers do not give child ratings that are relative to the school or class, so we are unable to do the same check as we did with parents.

in terms of parental effort or investment choices. The first step is decide which type of parental investment to analyze. In Section 2 we introduced two types of investment, "remedial" investment and "activities" investment. "Remedial" investment is calculated as the principal factor of the number of times per week that parents help their child with homework, the fraction of the times per week that parents help their child with homework, and whether the child is tutored. "Activities" investment follows the broad investment types discussed in Cunha *et al.* (2010). We use the principal factor of how often the parent: reads, tells stories, plays games, does science projects, dines together, goes to concerts, and goes to museums with child. We also include variables indicating whether the child takes music classes, art classes, drama classes, and sporting classes. Finally, we check whether these investment measures seem to respond to parental beliefs. The main results are reported in Table 9.

The first thing to notice about Table 9 is the timing. Parents are asked about their investment behavior during the last school year, while teacher ratings and the math score are relative to the spring term. For this reason, we treat parental investment as having been decided before the contemporaneous year information has been released. Hence, when we look at the variables that influence investment we used lagged test scores and lagged teacher ratings. The first four columns pertain to remedial investment, while the last four columns pertain to "activities" investments.

Consider remedial investment first. The first column indicates that parents who believe their child is above average invest 0.18 standard deviations less than other parents holding constant income, grades, and demographic variables such as race and gender. As shown in the previous section, we believe that beliefs are driven in part by math test scores and teacher ratings. We check the consistency of the previous result by controlling explicitly for test scores and teacher ratings. As expected, both measures negatively impact investment, although teacher reports are far more important (both measures are standardized). To understand whether investment seems to respond to change in beliefs we control in the third column for past variables. The coefficient on beliefs now indicates the impact on investment holding constant past beliefs and past investments. The coefficient is largely unchanged meaning that investment follows movements in beliefs. In the fourth column we control directly for the signals since they have the advantage of being continuous variables. Again the results do not change significantly after controlling for past investment and past signals.

In the next four columns we look at the "activities" definition of investment. This definition is strongly positively correlated with beliefs. This is more consistent with a story of dynamic complementarity where parents want to invest more in more able children. Once we control for past variables in the second to last column, we see that most of this effect is driven by ex-ante heterogeneity. Parents that start with positive beliefs tend to provide their children with more of this investment. This story is also supported by the fact that this type of investment is not strongly correlated, and does not strongly co-move, with observable measures of cognitive ability. While these investments might be very important for the cognitive development of a child, as demonstrated in Cunha *et al.* (2010), parents do not seem to be very elastic in their provision with respect to child scholastic performance.[9]. Given that we want to focus on investments that are potentially distorted by biased parental beliefs, in this paper we focus only on remedial investments as in the first four columns.

In Table 10 we test the robustness of our findings relating remedial investments to beliefs. One concern is that investment could be driven by different school policies. For example, in some schools children might be pushed harder and therefore require more help from their parents. This could be correlated with the average ability of the children, ex-ante and ex-post. To control for this, we incorporate school fixed effects into our regressions. While the point estimate is lower (-0.171), the difference is small and not statistically different from zero. In the third column, instead of including fixed effects, we control for the actual school policy. Teachers are asked how much time they expect their students to work on their homework each day. The coefficient on beliefs in this case increases, although not significantly.

[9]This idea is also consistent with the findings in table A10-2 of Cunha *et al.* (2010) where current period cognitive ability is not a significant determinant of parental investment, at least in the ages that overlap with our analysis

Another concern is whether parental investment responds to *global* beliefs or the child's *local* rank. This could be the case if, for example, parents care directly about their child's rank in the school or class. In the last two columns of Table 10 we try to address this concern. While we see that local relative ability does have an impact on investment, its presence does not reduce the impact of global beliefs on investment. Even in this case we might be concerned by the fact that our controls for beliefs are dichotomous. In the final column we control directly for continuous measures of cognitive ability. Conditional on the test score deviation from the school average, the level of the test directly impacts the level of investment. Note that in the last column we cannot interpret the math score as global beliefs and the test score deviation as local beliefs since the test score deviation impacts global beliefs. Yet, if the math score enters significantly, it strongly suggests that global beliefs enter significantly since the math score does not affect local beliefs.

3.4 Productive Effects of Investment

One implicit assumption we have made throughout the paper is that investment is productive. While we believe this is a reasonable assumption, it is not simple to identify in the data the causal impact of investment on achievement. In Table 11, we see that an OLS regression of math scores on past math scores and investment would produce a negative coefficient on investment. There is a clear endogeneity problem that we need to address: parents tend to give remedial inputs to children that are struggling. In the previous section we assumed that parental inputs are a function of past beliefs, and under this assumption we should not have any endogeneity concerns. However, in reality we cannot rule out that parents also respond with remedial inputs to contemporaneous signals of their child's productivity. The negative coefficient on investment is robust to the inclusion of demographics or school dummies.

In order to eliminate the endogeneity we need to utilize instrumental variables, i.e. we need to find a variable that is correlated with investment but uncorrelated with math scores conditional on past math scores. We use a factor based on measures of parental investment from the fall of kindergarten as our instrument. In our factor analysis we include variables like how often a parent reads, tells stories, sings songs, or play games with a child. Our assumption is that these activities are not correlated with math scores conditional on lagged math score, which is included in the regression. We also believe (and the first stage corroborates) that parental investment is correlated over time and that part of this correlation is the result of parental taste heterogeneity. Some parents are just more likely to spend time with their children regardless of their academic performance. When we apply the instrument, the coefficient becomes positive and significant. Controlling for demographics and school dummies reduces the magnitude a bit but does not change the general result.

An obvious concern with our instrument is that it is correlated with the test score residual though other types of investment. This would invalidate our instrument and bias the estimated impact of remedial type investments. As a result, we view the IV estimate as an upper bound for the impact of investment and the OLS estimate as a lower bound. Note that even if we target the negative reduced form effect of investment with our structural model it is still possible for investment to productive. Since test scores are noisy measures of ability, investment can also act as a (negative) proxy of past ability. This can lead to a negative coefficient even when investment is productive. Because of these myriad concerns, we will ultimately target different values for the productive impact of investment to get a sense for the robustness of our findings to this key parameter.

4 A Model of Beliefs, Information, and Investment

The evidence reported in the previous section suggests that parents have distorted beliefs about the relative ability of their children, these beliefs significantly influence investment, and that teachers are one potential channel through which parent's receive biased information. In this section we develop a structural model of parental beliefs, information, and investment that is compatible with the descriptive evidence. In following sections, we discuss estimation and the results from a series of counterfactuals.

The outline of the model is as follows. In each grade, a child is characterized by a (global) relative ability. This ability is unobserved both to parents and the econometrician. At the end of each grade, parents receive three signals about their child's relative ability: a test score, a teacher rating, and an additional signal that is unobserved to the econometrician. We assume that parents interpret each signal as an unbiased measure of a child's relative ability. However, we allow both the teacher signal and the unobserved signal to be distorted by the child's ability relative to the average ability in the child's school. This is motivated by our findings that parent and teacher beliefs are significantly driven by a local measure of scholastic performance. Parents use these signals to update their beliefs about ability and then choose investment in order to maximize the probability that their child will obtain a four-year college degree.

Fundamentals. In our model children begin in kindergarten ($g = 0$) and attend primary school for G periods. Children attend different schools (indexed by j) and are assumed to never change school. We model a child's relative ability, i.e. the ability of the child relative to the average ability in the economy. We assume that relative ability A_g is unobserved by both the econometrician and the parents. We suppress individual subscripts to minimize on notation. The average relative ability in school j is $\widetilde{A_{j,g}} = \frac{1}{N_j} \sum A_g$. Children are also diferentiated by a vector of observable (by both the econometrician and the parents) characteristics, X. Examples of variables included in X are family income, race, and gender, and the average of these variables in the child's school.

Signals. At the end of grade g, parents receive three signals about their child's ability: a test score S_g, a teacher report T_g, and a third "local" signal L_g. While the parent observes all of these signals, the econometrician does not observe the local signal, L_g. We assume that parents treat the three signals as unbiased but noisy measures of their child's relative ability. In other words, they act according to:

$$M_g = \gamma^M A_g + e_g^M \text{ for } M \in \{S, T, L\},$$

where each e_g^M is a mean zero independent random variable. A critical aspect of the model is that some of the signals are "locally biased;" that is, the process generating them is pulled toward the local average. While the test score S_g is unbiased, the teacher report and the local signal are both biased. We assume that the process generating these signals is:

$$M_g = \gamma^M \left(\alpha^M A_g + (1-\alpha^M)\left(A_g - \widetilde{A_{j,g}}\right)\right) + e_g^M \text{ for } M \in \{T, L\},$$

where the difference $A_g - \widetilde{A_{j,g}}$ measures a child's ability relative to his/her schoolmates. If $\alpha^M = 1$, the process generating measure M corresponds to the parent's interpretation of the signal. If $\alpha^M = 0$, parents believe that the signal is about their child's global relative ability when in reality it is entirely about their child's local relative ability. For any $0 < \alpha^M < 1$, the signal contains some good information and some systematic bias.

The complete parental information set in grade g is therefore $\Omega^g = (\{\Omega_n\}_{n=0}^g, X)$ where $\Omega_g = (S_g, T_g, L_g)$ contains the signals received in grade g. Notice that this information set is larger than the econometrician's, since he does not observe L_g.

The evolution of relative ability. Although children are not allowed to change school after entering kindergarten, we allow for initial sorting across schools on observables. We assume that the initial average relative ability in the school of a child is $\widetilde{A_{j,0}} \sim N\left(\mu(X), \sigma^2_{\widetilde{A}}\right)$; and that the initial ability of the child is drawn from: $A_0 \sim N\left(\widetilde{A_{j,0}}, \sigma^2_A\right)$.

The child's rank varies over time as a result of unpredictable shocks and parental investment:

$$A_g = A\left(A_{g-1}, I_g - \bar{I}_g, X\right) + u_g,$$

where I_g represents parental investment, $\bar{I}_g$ is the average investment in the whole economy and u_g an idiosyncratic mean zero shock. Notice that we subtract global average investment from parental investment. Given that the model is written in terms of *relative* ability, an increase in parental investment that is matched by an economy-wide increase in investment would not alter a child's global relative ability. Therefore, a child's g-grade relative ability is determined by his/her parents' relative investment, $I_g - \bar{I}_g$. In the estimation we assume that production takes the log-CES functional form:

$$A_g = \left(\frac{\gamma}{\rho}\right) \log\left(\pi \exp(\rho A_{g-1}) + (1-\pi)\exp(\rho(I_g - \bar{I}_g))\right) + u_g$$

For this function, the marginal productivity of investments is an increasing function of past ability if $\rho < 1$. This is referred to as "dynamic complementarity" by Cunha and Heckman (2010). On the other hand, if $\rho > 1$, the marginal productivity of investment is greater insofar as lagged ability is lower. In contrast to Cunha (2013), we assume that parents know this production function.

Parental utility and investment. We suppose each parent seeks to maximize an increasing function of her child's end-of-primary-school relative ability. Motivated by our data, which includes parents' reports regarding how far they expect their child will go in school, we make the plausible assumption that the driving goal of each parent is for her child to complete a four-year college degree. We then add the strong assumption that the probability of completing college depends only on child relative ability. This could be motivated by a short-run inelastic supply of college seats. We denote the probability of completing college as $C(A_G, X)$, where we allow, for example, the probability to be a function of variables like income or gender.

Parents choose their investment level in order to maximize their objective function. In order to make their choices, parents form expectations about the relevant variables. In particular they form their beliefs about A_g given their information at grade g, Ω^g. Denote the density function of these beliefs by $f(A_g|\Omega^g)$. An important feature of our data set is that we directly observe whether parents think their children are above average, both globally and locally, which we map into the model as the following objects: $\mathbb{1}\left(Pr\left(A_g|\Omega^g\right) > K_G\right)$ and $\mathbb{1}\left(Pr\left(A_g - \widetilde{A_{j,g}}|\Omega^g\right) > K_L\right)$, where K_G and K_L are constants. In words, if the chance her child is above average exceeds K_G, the parent says so; if the chance her child is above the school's average rank exceeds K_L (another constant), then she says so. Finally, as previously mentioned, we observe in grade G whether parents believe their children will complete college education. Analogously to the other reports, this maps into the model as $\mathbb{1}\left(Pr\left(C\left(A_G\right)|\Omega^G\right) > K_C\right)$, for a constant K_C.

Solving the model. Starting backwards, parents need to choose the optimal level of investment according to the following problem:

$$V_G(\Omega^{G-1}) = \underset{I_G}{max} \ \left[E\left(C(A_G, X)|\Omega^{G-1}\right) - \lambda_G(X) \times I_G^2\right].$$

We allow the marginal disutility of parental investment to be a function of X. While it is parsimonious to think of the state space in terms of Ω^{G-1}, note that parents only need $(f(A_g|\Omega^g), X)$ in order to make their optimal choice. Also, note that we are assuming that at the time of choosing the optimal investment in grade G, parents have only received signals up to period $G-1$. Therefore our model will not be able to capture patterns in investment that are generated by information that parents have received within the last year–say, additional information received in September that leads to additional tutoring in November.

The value function for earlier grades is similar:

$$V_g(\Omega^{g-1}) = \underset{I_g}{max} \ \left[\beta E\left(V_{g+1}(\cdot)|\Omega^{g-1}\right) - (\lambda_g(X) + e^{\lambda})I_g^2\right].$$

From the equation above, we see that in our model parents do not provide investment in kindergarten. Frankly, this is because we do not observe remedial investment choices in kindergarten. The initial period is still useful for estimation, as it provides information for the parents to construct the conditional distribution of their child's ability in first grade.

We assume that parents' do not directly observe other parents' investment decisions. This means that parents need to calculate $\bar{I}_g$ from their beliefs. Parents are assumed to know the distribution functions of all variables, so they compute aggregate ability by solving for:

$$\bar{I}_g = \int \left(argmax_I \int V_{g+1}(\cdot)dF\left(A_g|\Omega^{g-1}, I - \bar{I}_g\right) - \lambda(\cdot)I^2\right)dF\left(\Omega^{g-1}\right),$$

which under standard assumptions is a contraction mapping in $\bar{I}_g$. In solving the above problem, parents do not account for the fact that all the other parents have biased beliefs. As a result, parent calculations of aggregate investment will not match actual aggregate investment. Generally, parents will expect people to invest more than the actual amount since they use the "unbiased" distribution of ability to solve for $\bar{I}_g$.

Additions and extensions. In the next section we estimate the model as it has been presented thus far. However, we think it is useful to point out further extensions that we plan to incorporate. In the model, parents do not care about local relative ability. They are affected by it because the signals are biased but this measure of ability does not have a direct impact on parental choices or outcomes. We think this is a limitation for several reasons. School quality, measured by the average relative ability in the school, may influence the evolution of ability. Furthermore, local ability might influence college enrollment and eventually college graduation, as the secondary school GPA is normally utilized by colleges when considering students applications. In light of this discussion, we let the production function be:

$$A_g = A\left(A_{g-1}, \widetilde{A_{j,g-1}}, I_g - \bar{I}_g, X\right) + u_t,$$

while the probability of completing college would be $C(A_G, \widetilde{A_{j,G}}, X)$.

When parents care about the local distribution, they need to form expectations about it, i.e. they need to construct the joint density $f(A_g, \widetilde{A_{j,g-1}}|\Omega^g)$. In order to form a meaningful expectation of the school average, parents need to observe something related to it. One solution is to allow parents to observe an additional noisy signal (unobserved to the econometrician):

$$F_{j,g} = \gamma^F \widetilde{A_{j,g}} + e^F_{j,g}.$$

In addition, rather then assuming that parents perceive the teacher report and the local measure as unbiased measures of the global relative ability, we let these measures be an average of global and local ability, allowing the perceived weight to be different from the true one:

$$M_g = \gamma^M \left(\hat{\alpha}^M A_g + \left(1 - \hat{\alpha}^M\right)\left(A_g - \widetilde{A_{j,g}}\right)\right) + e^M_g \text{ for } M \in \{T, L\},$$

where in general $\hat{\alpha}^M \neq \alpha^M$ and we would expect that $\hat{\alpha}^M < \alpha^M$.

5 Estimation

We estimate our model using indirect inference. The model is targeted to match moments based on regressions similar to those described in Section 3: belief regressions, investment regressions, teacher rating regressions, and math scores regressions.

We first solve numerically for the value function, which delivers a mapping between beliefs about children's global relative ability and investment decisions. Even though we assume that all shocks are normally distributed, ability is in general not normally distributed as it evolves according to a non-linear law of motion as indicated by the log-CES production function. In order to make the problem tractable, we approximate beliefs using normal distributions and calculate how these distributions evolve using the Unscented Kalman Filter.[10]

We apply a few normalizations which are without loss of generality. Ability is never directly observed and therefore its scale is not identified (the mean is zero as this is relative ability). To set the scale we set the loading on ability for the test score $\gamma^T = 1$. We also set the loading on ability for the local signal $\gamma^L = 1$, as the measurement L_g is never directly observed.

Finally, we add measurement error in parental investment and assume that the production function is constant return to scale ($\gamma = 1$). The total number of model parameters is 22, although we will report only a subset of them.[11]

We utilize data for kindergarten, first and third grade. After obtaining a policy function, we simulate data for a large number of schools (10,000). We replicate the same distribution of pupils per school by randomly drawing from the data the number of students for each school. We store not only the simulated measures of rank, beliefs, and investment choices

[10] Cunha *et al.* (2010) utilize the same approach in a similar context.

[11] These parameters are: 3 variances that describe the distribution of children's ability, 8 parameters that describe the measurement equations for the signals, 2 parameters for the probability of attending college, 2 parameters for the production function, 2 parameters for the cost of investment, 1 parameter for identifying the threshold for being above average, 2 parameters for sample attrition, and 2 parameters of the measurement error of investment.

but also the underlying true values of child abilities which will be useful when evaluating counterfactuals. We mimic the structure of the data by also replicating sample attrition. After building the simulated data we compute a set of moments to compare to the ones produced by the actual data. The moments we use are:

1. An IV regression of math score on past math scores and investment, as discussed in Section 3.4;
2. An OLS regression of teacher rating on math scores and locally demeaned math scores;
3. The fraction of students that are considered above average by their parents;
4. OLS regressions of parental investments on lagged teacher ratings and math scores for each grade;
5. OLS regressions of parental investments on lagged parental beliefs for each grade;
6. An OLS regression of parental beliefs on teacher ratings and math scores;
7. An OLS regression of parental beliefs on math scores and locally demeaned math scores;
8. An OLS regression of parental beliefs regarding college completion on parental beliefs regarding children ability;
9. An OLS regression of parental beliefs regarding college completion on math score and teacher ratings;
10. Variances of test score and teacher ratings, covariance between math scores and school average math scores, variances of parental investment, covariance of teacher rating across different grades;
11. Probabilities of attrition.

The total number of coefficients that we match is 38. We estimated the model twice. The first estimation uses the optimal weighting matrix which is the inverse of the asymptotic covariance matrix of the above moments. The second estimation uses the identity matrix to weight the different coefficients. In Table 12, we report selected parameter estimates for the two cases. The estimates are generally similar, with a couple of important exemptions. The shock to the evolution of the child ability is much larger in the second case, as is the elasticity of substitution. The estimates of the parameters α indicate that both measures, teacher ratings and the local signal, are heavily biased towards the local distribution.

The elasticity of substitution indicates that the amount of dynamic complementarity is relatively low. Our two estimates are, for example, higher than those found in Cunha *et al.* (2010). Their first stage parameter, which is comparable in terms of the age of the child to our parameter, is 2.57 which is a bit smaller than our estimate using the optimal weighting matrix at 3.15. Our estimate when using the the identity matrix is instead much larger. The crucial reason for why we find lower dynamic complementarity is that we focus on remedial investment and, as shown in our data section, there is strong evidence that this investment is negatively related to ability.[12]

In Table 13, we report the main coefficients that we try to match in our estimation strategy. Although in general we can replicate fairly well the coefficients, a few remarks should be made. In general we fail to reproduce the positive impact of investment on future math scores. Using the optimal weighting matrix the failure is extreme as the coefficient on the simulated data is negative while it is positive, albeit smaller, when using the identity matrix. It is also worth pointing out that when using OLS on the actual data, the estimate of the coefficient for parental investment was negative and the positive coefficient was obtained using an IV estimation. Although this positive coefficient is statistically different from zero, it is not very precisely estimated and therefore the optimal weighting matrix

[12] In the data section we show that an alternative definition of investment, obtained using similar variables to Cunha *et al.* (2010) does not show the same negative relationship. The reason for which we focus on remedial investment is indeed its strong relationship to cognitive measures, which suggests that it is the most distorted investment type by the bias in beliefs.

does not give a lot of weight to it. Interestingly, even if the estimate is negative, the impact of investment on achievement is positive, as indicated by a positive π, the weight of investment on the production function, and by the results of the next section.

6 Counterfactual Exercises

Using the estimates from our structural model we investigate how investment and student test scores change when parents are better informed about their child's ability relative to children of the same age. There are two ways to achieve this, either remove the bias from the teacher and local signals or re-assign students so that there is no ability sorting into schools. Both approaches yield the same basic result.

Simulated data from our model has the added benefit that we can observe not only the observables but also the unobservables, i.e. when we simulate our data we store also the latent relative ability of the children. In the first panel of Tables 14 and 15 we simulate data using the structural estimates. Table 14 reports the results using the optimal weighting matrix and Table 15 shows the results using the identity matrix. In the first two columns we report the mean and standard deviations of the simulated third-grade ability and investment. We do so to remind the reader of the magnitude of these variables. In our simulation exercise we concentrate on the lowest part of the achievement distribution. The model indicates that these children are most affected by the bias in beliefs. Additionally these are the natural targets for educational policies. In the third column we report the average third-grade relative ability of the worst 10% of students as predicted by the model and their parental investments. In the fourth column we report the average relative ability for the students in the worst 10% of schools, where schools are ranked using the students average ability and their parental investments. In the second panel we look at how these numbers change when we remove the bias, namely we change α^T and α^L to one. In both the third and the fourth column the numbers relative to the average ability increase. The average ability for the worst 10% os students increases by 0.1 or 0.32 standard deviations. This increase is solely due to the increase in the parental remedial investments, which increase by around 0.2 to 1 standard deviation, depending on the weighting matrix and grade. The fourth column, which looks at the bottom 10% of schools, produces similar but slightly smaller impacts since the average ability in low achieving schools is not as low as in the previous case. Ability increases by 0.07 or 0.24 standard deviations, while investment increases by 0.2 to 0.7 standard deviations. Not shown in these tables, this exercise also implies that the top students perform a little worse than in the baseline specification, although the magnitude is not as large as for the bottom students (around one order of magnitude smaller).

One interesting piece of information that originates from our simulated data is reported in columns 5 and 6. Here we show the average ability and the average investment for the better half of students in the bottom 10% of schools and the worst half of students in the top 10% of schools. The better students in bad schools have much lower ability than bad students in good schools (-0.45 versus 0.34) and yet they receive less compensatory investment, as shown by the next two rows.

Table 16 shows a similar exercise seen from a different angle for both estimations. Instead of artificially removing the bias by setting the estimates equal to one, we manipulate the sorting into schools. Of course, changing the sorting would impact many other dimensions that are not included in our model. Yet we think this is a useful exercise to examine the implications of sorting from a different perspective. Each panel of Table 16 corresponds to a different weighting matrix. The first row of each panel reports third-grade ability while the following two rows report first-grade and third-grade investments. The first two columns are taken from the previous two tables and show the average quantity of those variables for the top 10% and bottom 10% of students. In the third and fourth column we show how these numbers change when we remove sorting. When we remove sorting, the local relative ability becomes in general much closer to the global relative ability. Local relative ability is not identical to global relative ability in this case either because of the small sample size in each school, and this explains why the numbers are different from the previous exercise. Relative ability for the bottom students increases by 0.1 and 0.27 standard deviations, while the decrease for the top students is much smaller in magnitude. Again, this change is completely driven by the change in the level of parental remedial investment, which decreases for top students and increases for bottom students. In the last two columns we look at the impact of increasing sorting in schools. In order to do this we look at the variance in within school initial ability (σ^2_A). We reduce that coefficient by 50% and increase the variance in the local average ($\sigma^2_{\tilde{A}}$) by the same amount. Relative ability drops by 0.08 or 0.20 standard deviations for the bottom students and barely increases for the top students as a result of changes in parental behavior.

Both exercises presented in this section point to the same result. In our model parents care about global ability but global ability is distorted toward the school relative ability. Removing this bias, either by giving parents the right information or by eliminating the difference between global and local relative ability, parents of low performing students realize the extent of their child's scholastic achievements and react by increasing the level of their remedial investment, namely helping their children with homework or offering them tutoring.

7 Concluding Remarks

In this paper we present evidence that parental beliefs about a child's ability relative to similarly aged children are distorted by a child's ability relative to children in the same school. This distortion in beliefs has important consequences for parental investment and the evolution of children's ability. Parents of low ability children who attend schools where average ability is also low will perform fewer remedial type investments than parents of similarly able children who attend schools where average ability is higher. Because of the tendency for students and families to sort into schools and neighborhoods, low ability children are more likely to attend schools where average ability is also low. As a result, the distortion in parental beliefs generated by local ability comparisons leads to underinvestment for low ability children. Using our structural model, we find that the impact of distorted parental beliefs on achievement can be quite large at the bottom of the ability distribution.

Our paper complements recent work illustrating parental misinformation about child ability and development. Cunha *et al.* (2013) finds that socioeconomically disadvantaged, pregnant African American women have biased beliefs regarding the productivity of parental investment. Closer to our paper, Dizon-Ross (2013) finds that parents in Malawi significantly overstate their child's ability and when given more accurate information choose more remedial type investments to help their children. The elicitation of parental beliefs in these papers is cleaner than in our setup since the survey/experiments employed were designed precisely for this reason. However, by using the ECLS-K we are able to explore in greater detail the nature and source of parental distortions. As a result we are able to gain additional insight into policies capable of ameliorating these distortions.

The finding that parent beliefs about a child's relative ability are distorted by the local ability distribution also connects our paper to the broader peer effects literature. In an effort to estimate the impact of peers, researchers often estimate the impact average classroom ability has on individual test score outcomes. It is generally not clear the channel through which average peer ability operates, but the typical interpretation is that it works through in-school behaviors of either the teacher or students themselves. Our paper suggests that average peer ability also matters for individual outcomes through its impact on parental investment.

Parental investment in children, particularly at young ages, has been shown to a be a key input into skill development. As a result, it is imperative that we understand the key determinants of these investment decisions. Our paper suggests that one important factor are parental beliefs about the cognitive ability of their child. However, significant work remains. In particular, parental beliefs about the returns to investment and beliefs about non-cognitive skills are likely to significantly influence investment decisions. Moreover, embedding parental beliefs into a broader model of investment that accounts for borrowing constraints and the trade-offs between goods and time investments would be extremely informative. These additional constraints may temper the impact of beliefs or exacerbate them depending on the relationships between beliefs and other family characteristics.

References

Antonovics, Kate, & Golan, Limor. 2012. Experimentation and Job Choice. *Journal of Labor Economics*, **30**(2), 333–366.

Arcidiacono, Peter, Aucejo, Esteban, Maurel, Arnaud, & Ransom, Tyler. 2014. *College Attrition and the Dynamics of Information Revelation*. Working Paper.

Becker, Gary S., & Tomes, Nigel. 1986. Human Capital and the Rise and Fall of Families. *Journal of Labor Economics*, **4**(3), S1–S39.

Bernal, Raquel. 2008. The Effect of Maternal Employment and Child Care on Children's Cognitive Development. *International Economics Review*, **49**(4), 1173–1209.

Boca, Daniela Del, Flinn, Christopher, & Wiswall, Matthew. 2014. Household Choices and Child Development. *Review of Economic Studies*, **81**(1), 137–185.

Carneiro, Pedro, & Heckman, James J. 2003. Human Capital Policy. *In:* Heckman, James J., Krueger, Alan B., & Friedman, Benjamin M. (eds), *Inequality in America: What Role for Human Capital Policies*. Cambridge, MA: MIT Press.

Caucutt, Elizabeth M., & Lochner, Lance. 2012. *Early and Late Human Capital Investments, Borrowing Constraints, and the Family*. NBER Working Paper 18493.

Cunha, Flavio. 2013. *Investments in Children When Markets are Incomplete*. Working Paper.

Cunha, Flavio, & Heckman, James J. 2008. Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation. *Journal of Human Resources*, **43**(4), 738–782.

Cunha, Flavio, Heckman, James J., & Schennach, Susanne M. 2010. Estimating the Technology of Cognitive and Noncognitive Skill Formation. *Econometrica*, **78**(3), 883–931.

Cunha, Flavio, Elo, Irma, & Culhane, Jennifer. 2013. *Eliciting Maternal Expectations about the Technology of Cognitive Skill Formation*. NBER Working Paper 19144.

Dizon-Ross, Rebecca. 2013. *Parents' Perceptions and Children's Education: Experimental Evidence from Malawi*. Working Paper.

Heckman, James J., & Mosso, Stefano. 2014. The Economics of Human Development and Social Mobility. *Annual Review of Economics*, **6**, 689–733.

Kinsler, Josh, & Pavan, Ronni. 2014. *The Specificity of General Human Capital: Evidence from College Major Choice*. forthcoming Journal of Labor Economics.

Sanders, Carl. 2014. *Skill Accumulation, Skill Uncertainty, and Occupational Choice*. Working Paper.

Stinebrickner, Todd R., & Stinebrickner, Ralph. 2009. *Learning about Academic Ability and the College Drop-Out Decision*. NBER Working Paper 14810.

Todd, Petra E., & Wolpin, Kenneth I. 2007. The Production of Cognitive Achievement in Children: Home, School, and Racial Test Score Gaps. *Journal of Human Capital*, **1**(1), 91–136.

Wiswall, Matthew, & Zafar, Basit. 2014. *Determinants of College Major Choice: Identification using an Information Experiment*. forthcoming Review of Economic Studies.

Figure 1: Fraction of Parents Believing in Child College Graduation

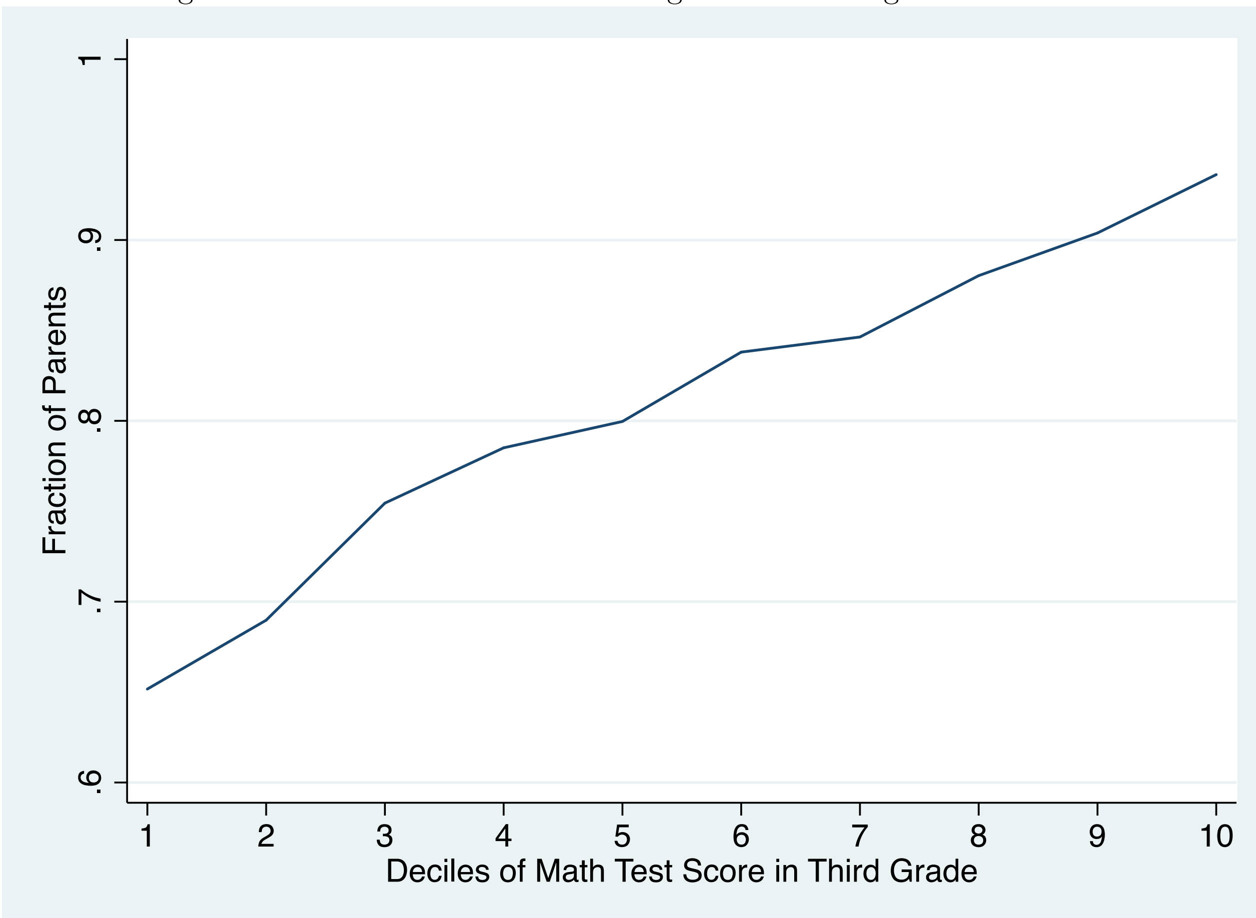

Table 1: Summary Statistics

	K	1st	3rd
% White	0.55	0.57	0.60
Log Income	10.47	10.52	10.57
Children per School	21.6	16.5	12.9
Comparisons to children of same age			
Above Average	0.34	0.31	0.34
Below Average	0.07	0.07	0.08
Comparisons to children in same class			
Above Average, Math		0.36	0.35
Below Average, Math		0.05	0.07
Above Average, Reading		0.40	0.36
Below Average, Reading		0.09	0.09
Parents believe child will obtain 4-yr degree	0.79	0.76	0.81
Parents help with HW, 5+ times per week		0.28	0.22
Parents help with HW, 1-2 times per week		0.21	0.30
Parents help with HW, Never		0.05	0.06
N	19,576	14,940	11,253

Table 2: Parental Beliefs and Math Scores

	Above Average Relative to ...			
	Similarly Aged Children		Children in Same Class	
			Math	Reading
Math	0.158	0.107	0.019	0.003
	(0.003)	(0.006)	(0.007)	(0.007)
Math - School Avg. Math		0.069	0.136	0.138
		(0.007)	(0.007)	(0.008)
Grade Controls	Y	Y	Y	Y
N	23,048	23,048	23,090	23,104

Table 3: Parental Beliefs and Reading Scores

	Above Average Relative to ...			
	Similarly Aged Children		Children in Same Class	
			Math	Reading
Reading	0.162	0.112	0.012	0.023
	(0.003)	(0.006)	(0.007)	(0.006)
Reading - School Avg. Reading		0.068	0.102	0.176
		(0.007)	(0.008)	(0.007)
Grade Controls	Y	Y	Y	Y
N	22,775	22,775	22,816	22,830

Table 4: Parental Beliefs, Robustness to ME

	Above Avg. Relative to Similarly Aged Children				Read for Fun ≥ 3 times per week	
	OLS	IV	OLS	IV	OLS	OLS
Math	0.107	0.125			0.095	
	(0.006)	(0.007)			(0.006)	
Math - School Avg. Math	0.069	0.117			-0.047	
	(0.007)	(0.009)			(0.006)	
Reading			0.112	0.138		0.105
			(0.006)	(0.007)		(0.005)
Reading - School Avg. Reading			0.068	0.101		-0.024
			(0.007)	(0.009)		(0.006)
Grade	Y	Y	Y	Y	Y	Y
N	23,048	22,812	22,775	22,170	23,170	22,897

Table 5: Parental Beliefs, Robustness to Varying Reference Points

	Above Average Relative to Similarly Aged Children					
	using math score controls			*using reading score controls*		
Test Score	0.107	0.109	0.080	0.112	0.128	0.124
	(0.006)	(0.009)	(0.021)	(0.006)	(0.009)	(0.023)
School Deviation	0.069	0.070	0.067	0.068	0.074	0.069
	(0.006)	(0.008)	(0.07)	(0.007)	(0.008)	(0.007)
Socioeconomic Deviation		-0.003			-0.023	
		(0.010)			(0.010)	
Geographic Deviation			0.029			-0.013
			(0.022)			(0.024)
Grade	Y	Y	Y	Y	Y	Y
N	23,048	23,048	23,048	22,775	22,775	22,775

Table 6: Parental Beliefs, Robustness to School Heterogeneity

	Above Average Relative to Similarly Aged Children							
Math	0.107	0.112			0.130	0.148		
	(0.006)	(0.007)			(0.004)	(0.006)		
Math - School Avg.	0.069	0.068						
	(0.007)	(0.008)						
Math - Class Avg.					0.056	0.038		
					(0.006)	(0.007)		
Reading			0.112	0.121			0.136	0.154
			(0.007)	(0.007)			(0.004)	(0.006)
Reading - School Avg.			0.068	0.062				
			(0.007)	(0.008)				
Reading - Class Avg							0.050	0.032
							(0.006)	(0.007)
Assigned HW	N	Y	N	Y	N	N	N	N
School FE	N	N	N	N	N	Y	N	Y
N	23,048	22,812	22,775	22,170	23,170	22,897	23,170	22,897

Table 7: Teacher Ratings

	Math Skills			Reading Skills		
Test Score	0.615	0.417	0.228	0.699	0.440	0.281
	(0.006)	(0.011)	(0.018)	(0.005)	(0.010)	(0.015)
Deviation from School Avg.		0.268	0.152		0.354	0.236
		(0.012)	(0.019)		(0.011)	(0.017)
Past Teacher Rating			0.251			0.307
			(0.007)			(0.007)
Lag Test Score			0.103			0.050
			(0.018)			(0.016)
Lag Deviation from School Avg.			0.102			0.052
			(0.020)			(0.018)
Grade Controls	Y	Y	Y	Y	Y	Y
N	22,855	22,855	21,432	22,645	22,645	20,984

Table 8: Parent Beliefs as a Function of Teacher Ratings

	Above Average Relative to Similarly Aged Children					
Math	0.158	0.092	0.060	0.058	0.040	0.053
	(0.003)	(0.004)	(0.004)	(0.004)	(0.007	(0.004)
Teacher Math Rating		0.105	0.094	0.038	0.043	0.051
		(0.004)	(0.004)	(0.005)	(0.009)	(0.006)
Beliefs, Comparison to Class	N	N	Y	N	Y	N
Lagged Controls	N	N	N	Y	Y	Y
School FE	N	N	N	N	N	Y
Grade Controls	Y	Y	Y	Y	N	Y
N	23,048	20,547	20,387	17,779	7,364	17,779

Table 9: Parental Investment

	Remedial				Activities			
Lag Beliefs	-0.180		-0.170		0.164		0.038	
	(0.015)		(0.025)		(0.015)		(0.022)	
Lag Math		-0.064		-0.065		0.002		-0.015
		(0.009)		(0.014)		(0.009)		(0.013)
Lag Teacher		-0.134		-0.124		0.025		0.005
		(0.009)		(0.017)		(0.009)		(0.0015)
Lag Investment			0.167	0.152			0.498	0.502
			(0.012)	(0.012)			(0.010)	(0.010)
Lag^2 Beliefs	N	N	N	Y	N	N	N	N
Lag^2 T and M	N	N	N	N	N	N	N	Y
Grade Controls	Y	Y	Y	Y	Y	Y	Y	Y
Demographics	Y	Y	Y	Y	Y	Y	Y	Y
N	21,150	21,386	8,364	8,214	21,022	21,247	8,271	8,117

Table 10: Parental Remedial Investment, Robustness

Lag Beliefs, Similarly Aged Chlidren	-0.180	-0.171	-0.192	-0.194	
	(0.015)	(0.015)	(0.016)	(0.025)	
Lag Beliefs, Class Comparison				-0.162	
				(0.025)	
Lag Math					-0.135
					(0.017)
Lag Math Deviation					-0.055
					(0.019)
School FE	N	Y	N	N	N
HW Policy	N	N	Y	Y	Y
Grade and Demographics	Y	Y	Y	Y	Y
N	21,150	21,150	18,501	7,920	19,797

Table 11: Productive Effect of Investment

	Math Score					
	OLS	IV	OLS	IV	OLS	IV
Lag Math Score	0.762	0.794	0.720	0.741	0.696	0.711
	(0.005)	(0.009)	(0.005)	(0.010)	(0.006)	(0.011)
Remedial Investment	-0.040	0.179	-0.041	0.098	-0.044	0.071
	(0.004)	(0.052)	(0.004)	(0.053)	(0.004)	(0.057)
School Dummies	N	N	N	N	Y	Y
Demographics	N	N	Y	Y	Y	Y
Grade Controls	Y	Y	Y	Y	Y	Y
N	22,919	20,233	22,413	20,127	22,413	20,127

Table 12: Key Parameter Estimates

	Description	Optimal	Identity
$\sigma_{\tilde{A}}$	Local Mean Ability	0.446	0.493
σ_A	Local Ability	0.744	0.603
α^T	1 - Bias in Teacher Signal	0.393	0.407
α^L	1 - Bias in Unobserved Signal	0.124	0.182
$\sigma_{e_g^S}$	Test Score	0.670	0.583
$\sigma_{e_g^T}$	Teacher Signal	0.537	0.413
$\sigma_{e_g^L}$	Unobserved Signal	0.117	0.210
$\frac{1}{1-\rho}$	Elasticity of Substitution	3.152	41.820
$1-\pi$	Weight on Investment	0.007	0.009
σ_{u_g}	Production Shock	0.033	0.610
$\sigma_{\epsilon_I^1}$	ME Investment	0.996	0.867
$\sigma_{\epsilon_I^3}$	ME Investment	0.886	0.715

Table 13: Moments Fit

Dep. Variable	Regressor		Weighting Matrix	
		Data	Optimal	Identity
Cognitive Measures				
Math Score	Lag Score	0.794	0.706	0.758
	Investment	0.179	-0.013	0.109
Teacher Signal	Test Score	0.395	0.432	0.449
	Local Math Rank	0.279	0.230	0.326
Parental Investment				
Investment - 1st g.	Above Average	-0.140	-0.140	-0.131
Investment - 3rd g.	Above Average	-0.243	-0.180	-0.230
Investment - 1st g.	Teacher Score	-0.047	-0.049	-0.083
	Test Score	-0.098	-0.052	-0.115
Investment - 3rd g.	Teacher Score	-0.088	-0.075	-0.120
	Test Score	-0.156	-0.096	-0.169
Beliefs				
Above Average	Teacher Signal	0.092	0.145	0.136
	Test Score	0.105	0.160	0.192

Table 14: Eliminating Parental Biases - Optimal Weighting Matrix

	Summary St.		Bottom 10%		School Distribution	
	Mean	SD	Students	Schools	< 10% Above LA	> 90% Below LA
Baseline						
Ability, 3rd	0.000	0.839	-1.422	-0.771	-0.210	0.180
Invest, 1st	0.118	1.012	0.349	0.131	0.042	0.158
Invest, 3rd	0.146	0.984	0.653	0.179	0.056	0.196
$\alpha^{\mathbf{T}} = \mathbf{1}$ **and** $\alpha^{\mathbf{L}} = \mathbf{1}$						
Ability, 3rd			-1.337	-0.711		
			(0.10 SD)	*(0.07 SD)*		
Invest, 1st			0.546	0.327		
Invest, 3rd			1.168	0.587		

Table 15: Eliminating Parental Biases - Identity Weighting Matrix

	Summary St.		Bottom 10%		School Distribution	
	Mean	SD	Students	Schools	< 10% Above LA	> 90% Below LA
Baseline						
Ability, 3rd	0.000	1.090	-0.940	-0.754	-0.452	0.344
Invest, 1st	0.125	0.992	0.633	0.225	0.048	0.099
Invest, 3rd	0.171	0.994	0.706	0.292	0.081	0.151
$\alpha^{\mathbf{T}} = \mathbf{1}$ **and** $\alpha^{\mathbf{L}} = \mathbf{1}$						
Ability, 3rd			-0.596	-0.491		
			(0.32 SD)	*(0.24 SD)*		
Invest, 1st			1.590	0.943		
Invest, 3rd			0.947	0.538		

Table 16: Role of Sorting

	Student Distribution					
	Baseline		No Sorting		More Sorting	
	< 10%	> 90%	< 10%	> 90%	< 10%	> 90%
Optimal						
A, 3rd	-1.422	1.496	-1.339	1.488	-1.491	1.498
			(0.10 SD)	*(-0.01 SD)*	*(-0.08 SD)*	*(0.002 SD)*
I, 1st	0.349	0.023	0.551	0.006	0.205	0.041
I, 3rd	0.653	0.008	1.133	-0.004	0.229	0.026
Identity						
A, 3rd	-0.940	1.279	-0.606	1.237	-1.162	1.303
			(0.27 SD)	*(-0.04 SD)*	*(-0.20 SD)*	*(0.02 SD)*
I, 1st	0.633	0.027	1.571	0.012	0.215	0.033
I, 3rd	0.706	0.009	0.932	0.019	0.457	0.021