Parents' Perceptions and Children's Education: Experimental Evidence from Malawi

Rebecca Dizon-Ross*

JOB MARKET PAPER

November 14, 2014

Latest version: http://www.mit.edu/~rdr/perceptions.pdf

Abstract

Do parents' inaccurate beliefs about their children's academic skills distort investments in education? I analyze a field experiment conducted in Malawi and find that, at baseline, parents try to tailor their educational investments to their children's academic level, but their inaccurate beliefs prevent them from doing so. Providing parents with information about their children's academic performance causes them to update their beliefs and align their investments more closely with their children's achievement. For example, most parents think that paying for secondary school is more valuable for higher-achieving children, but some parents are mistaken about their children's relative achievement; parents who receive information reallocate resources towards their higher achievers. Poorer, less-educated parents have less accurate beliefs about their children than richer, more-educated parents, and update their beliefs and investments more in response to information. Inaccurate perceptions may thus exacerbate inequalities between richer and poorer households or societies.

* I am very grateful to Pascaline Dupas, Caroline Hoxby, and Seema Jayachandran for their guidance throughout all stages of the project, and to Ran Abramitzky, Abhijit Banerjee, Jim Berry, Nick Bloom, Doug Bernheim, Celine Dizon, Elise Dizon-Ross, Natalie Douvos, Esther Duflo, Alex Eble, Liran Einav, Nick Hagerty, Rema Hanna, Johannes Haushofer, Anil Jain, Asim Khwaja, Anjini Kochar, Shirlee Lichtman, Matthew Lowe, Rachael Meager, Ben Olken, Arianna Ornaghi, Jonah Rockoff, Sheldon Ross, Ashish Shenoy, Fabiana Silva, Melanie Wasserman, and Jenny Ying for helpful discussions that greatly benefited the project. I thank participants at the Columbia Teacher's College, Harvard Development Lunch, MIT Development Lunch, NEUDC 2013, NBER Development Fall 2014, NBER Education Spring 2014, PacDEV 2014, Stanford Applied Lunch, and Stanford Development Seminar for helpful comments. I also thank Bridget Hoffmann, Rachel Levenson, and Michael Roscitt for help with the fieldwork. I appreciate the generous support of the endowment in memory of B.F. Haley and E.S. Shaw, Innovations for Poverty Action, the National Science Foundation (DDRIG 1156155), the Russell Sage Foundation, the Shultz Graduate Student Fellowship in Economic Policy, the Stanford Economics Department, the Stanford Institute for Economic Policy Research, and the Stanford DDRO and GRO Funds. All errors are my own. Contact: Abdul Latif Jameel Poverty Action Lab, MIT, E53-389, 30 Wadsworth St., Cambridge, MA 02142. rdr@mit.edu.

1 Introduction

It is widely believed that efficient investments in human capital vary based on individual characteristics. This idea is embeddded in a long line of models for optimal human capital investment, dating back to Becker (1962), in which the returns to education depend on some individual characteristic, like academic ability.

However, it is an open question whether actual schooling investments vary efficiently with individual characteristics. In particular, since it is perceived, not true, characteristics that govern educational investments, inaccuracies in parents' perceptions about their children's characteristics could cause distortions. To give a concrete example, consider a parent with two children, one who performs well in school and one who does not. The parent can only afford to send one child to secondary school, and wants to send her higher-achieving child. However, she has inaccurate beliefs about which of her children is higher-performing, and so chooses to send the lower achiever, only to have that child fail out of secondary school.

Misallocations caused by inaccurate perceptions could be particularly problematic in developing countries because many parents are uneducated. Access to primary schooling only became free in many countries in the last 10-20 years: the average adult in sub-Saharan Africa has fewer than 5 years of education (UNESCO, 2013). Limited education and illiteracy may make it difficult for parents to judge their children's academic performance, especially if their children go further in school than they did, as is common in developing countries. Banerjee et al. (2010) find that, in India, 55% of parents whose child can barely decipher letters think the child can read paragraphs and stories. Combining data from my sample in Malawi with data from the U.S., I see that, both within and across countries, less educated parents appear to have less accurate beliefs (see Fig. 1).[1]

As a result, parents' inaccurate perceptions could help explain why educational outcomes in developing countries are both poor and unequal. For example, in Malawi, the secondary school completion rate is below 50%, and is over twice as high for the richest quintile of households as for the poorest. For primary school, the age group I focus on in this paper, the completion rate is only 35%, and, despite primary school's lower costs, the differences between rich and poor even more stark (World Bank, 2007). Researchers have examined many factors (e.g., limited access to credit, school quality) to explain these patterns, but none fully account for them. The hypothesis examined here is that inaccuracies in uneducated parents' perceptions distort their investments, providing one channel for the perpetuation of inequalities across generations.

[1] The cross-country relationship should be interpreted with caution as the sampling methods and academic performance measures vary across countries. The U.S. data come from Alexander and Entwisle (2006).

This paper analyzes a field experiment conducted in Malawi that targeted parents with children in grades 2-6. I first measured parents' beliefs about their children's academic performance.[2] I then provided randomly selected parents with information about their children's true academic performance (specifically, their average scores and within-class percentile rankings on achievement tests administered in schools). Finally, I measured the effects on parents' educational investments. The information delivered was very similar to the information already nominally given to parents through report cards in many countries, including Malawi. However, the official report cards are often hard for parents to understand, or do not reach them. The intervention presented information more clearly.

In addition to identifying distortions due to inaccurate beliefs, this experiment allows me to uncover how parents' investments depend on their children's characteristics. Many papers have examined this relationship (e.g., Behrman et al. (1994)), but it is difficult to estimate, most notably because of the potential for reverse causality between investments and child characteristics. The experiment enables a new within-person methodology that exploits the exogenous "shock" to parents' beliefs for identification.

The experimental results imply that inaccuracies in parents' perceptions about their children's characteristics may have large, negative impacts on children's education. First, I find that parents' beliefs about their children's achievement diverge substantially from their children's true achievement; the gap is more than one standard deviation of the achievement distribution on average. Comparing two of their children, one third of parents are mistaken about which child has higher achievement. Belief inaccuracies also vary with parents' education, with less-educated parents having significantly less accurate beliefs.

Second, I establish that inaccurate beliefs distort investments by using lab-in-the-field style investments chosen to have clear predictions for how the efficient investment depends on achievement. For example, I study willingness-to-pay (WTP) for remedial math and English textbooks, which should decrease with achievement because the textbooks are remedial (i.e., perceived substitutes with achievement). In the control group, parents spend more on the book in the subject they think their child is behind in, but parents' beliefs are often inaccurate. Providing information causes parents to reallocate towards the subject in which their child is truly behind. I find similar results for parents' choices among free workbooks targeted for different achievement levels. In addition, there are larger distortions among less-educated parents: because they have less accurate beliefs, they update their beliefs and investments more in response to information.

Next, I test whether these distortions are relevant for parents' broader investments in education (i.e., the investments that we care most about, such as secondary schooling and

[2] I discuss the choice to focus on achievement, rather than "innate" ability, later in the paper.

dropouts). I find that they are. Secondary school fees are the first high-cost investment that parents in Malawi make in their children's education, and many parents cannot afford the fees for all of their children. To gauge whether parents misallocate these investments, I gave parents the chance to win a secondary school scholarship (more specifically, for budgetary reasons, entry into a lottery for a scholarship) and asked them to choose which of their children to give it to. Information causes parents to reallocate the scholarship towards the child that is higher-achieving. I find similar effects for dropouts: children whose parents found out that they were doing well in school are less likely to drop out, whereas children whose parents found out they were doing poorly are more likely to drop out. This behavior suggests that parents believe years of schooling and achievement are complements, which is consistent with parents' reported beliefs as well as the findings of the literature (Pitt et al., 1990; Aizer and Cunha, 2012).

One might expect that the patterns would vary for other types of investments, like expenditures: conditional on staying in school, parents may need to spend more on their weaker students to keep them on track. And in fact, when I examine expenditures, I find evidence that poorer parents spend more on their lower achievers. However, richer parents do the opposite. One potential explanation is that richer parents think they can send their children to secondary school and poorer parents do not, so richer parents want to get their high achievers over the admission threshold, whereas poorer parents think there are higher returns to helping low achievers acquire basic skills like learning to read. But, for both groups, information allows parents to reallocate.

Taken together, these results suggest that parents use the new information to try to optimize their investments. Whether reoptimization leads to an increase in efficiency depends on whether parents are right about the production function; this is an open question, but it is reassuring that parents' behavior is generally in line with our *ex ante* predictions. These results also highlight that information is not a universal panacea to increase education for all – it leads to reallocations, which can decrease investments for some groups.

These results help advance our understanding of the causes of poor educational outcomes in developing countries by showing that inaccurate perceptions about achievement may prevent some households from taking full advantage of educational opportunities. This paper also contributes to several strands of the literature. First, it contributes to a recent and growing literature on the role of information constraints in education. The majority of the literature focuses on misinformation about population-average characteristics, like the returns to education or information about school quality (Jensen, 2010; Nguyen, 2008; Dinkelman and Martínez A, 2014; Andrabi et al., 2009). There is some suggestive evidence that misinformation about individual-level characteristics causes misallocations, but, to my

knowledge, this paper performs the first causal test (Chevalier et al., 2009; Stinebrickner and Stinebrickner, 2009; Andrabi et al., 2009; Bergman, 2014).[3] More broadly, the previous literature does not disentangle the mechanisms for information's effects, such as whether information changes the point estimate or the uncertainty of perceived returns. By incorporating data on baseline beliefs, this paper can help disentangle the mechanisms and trace out the causal chain from beliefs to investments. Second, it contributes to the literature examining how parents' investments depend on their children's ability by using a new method for identification (e.g., Behrman et al., 1994; Griliches, 1979; Datar et al., 2010; Leight, 2014; Rosenzweig and Zhang, 2009). The literature's findings are mixed between compensation and reinforcement. This paper suggests that, even within a given context, parents' objectives vary across parents (e.g., rich vs. poor) and investments (e.g., dropout vs. spending on children in school), which could help explain the mixed results from the earlier literature. Finally, it adds to the literature documenting the positive influence of parents' education on children's education by highlighting a specific channel for effects: parents' beliefs (Rosenzweig and Wolpin, 1994; Oreopoulos et al., 2006; Andrabi et al., 2012; Banerji et al., 2013).

The remainder of the paper proceeds as follows. Section 2 presents the theoretical predictions. Section 3 describes the context and experimental design. In Section 4, I use the baseline data to examine whether parents have inaccurate perceptions and how that impacts their investments. Section 5 presents the results on the impact of information on the lab-in-the-field investments. Section 6 examines heterogeneity by parent education. Section 7 looks at longer-run outcomes. In Section 8, I conclude.

2 Theoretical predictions

This section presents a simple model to generate predictions for how inaccurate beliefs distort investments that I then test in the empirical analysis. In the model, parents have a fixed budget for education and one child. They allocate spending across two inputs, one that is a complement with achievement and one that is a substitute. This fixed-budget, one-child model is used for expositional simplicity, and to illustrate how inaccurate beliefs could cause misallocations conditional on the level of spending. However, the same theoretical predictions would hold under many other models, for example, if the education budget were flexible and

[3]The existing evidence mostly relies on correlations; for example, Chevalier et al. (2009) find that, conditional on true academic ability, beliefs about academic ability are correlated with educational outcomes. The two previous studies that provide exogenous variation in beliefs about individual characteristics are Andrabi et al. (2009) and Bergman (2014). Andrabi et al. (2009) bundles the individual-level information with school-level information so the intent is not to separately identify its effects. Bergman (2014) provides individual-level academic information to parents in a developed country context (Los Angeles in the U.S.), but the focus is on a different type of information inefficiency – agency issues – rather than misallocations of investments by individual characteristics.

all inputs complements. The same model could also be used to describe misallocations across children within the household.[4]

The model

A parent has a fixed budget for education to allocate between two inputs to her child's education. Her returns-maximizing allocation solves the following maximization problem:

$$\max_{s^c, s^s} \; q(s^c, s^s|A) \text{ s.t. } s^c + s^s \leq y^{educ}$$

where q is the education production function measuring the long-run output of the investment (e.g., earnings); A is the child's achievement level; s^c, the first input, is a complement with achievement $\left(\frac{\partial^2 q}{\partial s^c \partial A} > 0\right)$; s^s, the second input, is a substitute with achievement $\left(\frac{\partial^2 q}{\partial s^s \partial A} < 0\right)$; output is increasing in both inputs $\left(\frac{\partial q}{\partial s^c} > 0, \frac{\partial q}{\partial s^s} > 0\right)$; and y^{educ} is the total budget for education. For example, s^s could be a remedial tutoring class and s^c could be an advanced tutoring class. Note that achievement is used here as a (potentially noisy) measure of ability or, more precisely, individual-level returns – it is an input, not an output, to the static production function. I assume in the model that parents know the true production function, q, but discuss the implications if parents do not know the true function later in the paper.[5]

Under standard assumptions (e.g., concave returns), this problem yields unique, returns-maximizing choice functions for educational inputs, $s^{c*}(A)$ and $s^{s*}(A)$, with the complement increasing in achievement, $\frac{\partial s^{c*}}{\partial A} > 0$, and the substitute decreasing, $\frac{\partial s^{s*}}{\partial A} < 0$. For simplicity, I parametrize these relationships as $s^{c*}(A) = \beta_0^c + \beta_1^c A$ and $s^{s*}(A) = \beta_0^s + \beta_1^s A$, with $\beta_1^c > 0, \beta_1^s < 0$. To simplify notation, I hereafter refer to both investments as $s^*(A)$ and both slopes as β_1.

Assume that parents do not know their children's true achievement, A; instead, they have beliefs about achievement, $\tilde{A}$. For now, I assume there is no uncertainty in beliefs, but discuss uncertainty at the end of the section. Instead of choosing the returns-maximizing investments, $s^*(A) = \beta_0 + \beta_1 A$, parents instead choose investments that depend on their beliefs, $s^*(\tilde{A}) = \beta_0 + \beta_1 \tilde{A}$. As a result, if beliefs are inaccurate, they earn inefficiently low returns: $q(s^*(\tilde{A})|A) \leq q(s^*(A)|A)$. Note that I use the term "inaccurate beliefs" to characterize beliefs that diverge from true achievement at the individual level: $Abs.Val.(\tilde{A}_i -$

[4]To see this, simply relabel the two inputs as "spending on child 1" and "spending on child2" and take A to be the gap in achievement between child 1 and child 2.

[5]In general, revealed preference tests only teach us about perceived returns, not true returns; as such, by observing parents' behavior in this setting, one can test whether increasing beliefs accuracy improves output as measured by parents' *perceived* production function, or whether a given input is a *perceived* complement or substitute with achievement. This is one reason it is useful to study investments where we have clear hypotheses about the true production function, like remedial textbooks, and can test whether parents' perceived production functions conform to our *ex ante* expectations.

$A_i) > 0$. Inaccurate beliefs could either be unbiased ($E(\tilde{A}) = E(A)$) or biased ($E(\tilde{A}) \neq E(A)$). Both types of inaccuracies distort individual-level inputs, and thus average returns.

Testing for distortions

Inaccurate beliefs decrease returns by distorting parents' actual investments away from their desired investment functions, where I use the term "desired investment function" to represent how parents want to invest (that is, the $s^*(.)$ function which maps achievement to input choices to correctly solve parents' optimization function) and the term "actual investment function" to represent how parents actually invest (that is, the true (average) mapping from children's achievement to parents' input choices, $s(A)$). Gauging whether there are divergences between the two functions can thus help us evaluate whether inaccurate beliefs distort returns.

Figure 2 demonstrates how to do this. To determine the desired investment function, one can look at how investments depend on perceived achievement (the dotted lines). These lines should be steeply sloped with slope β_1. To determine the actual investment function, one can see how investments depend on true achievement (the solid lines: the line has the same y-axis as the dotted line but a different x-axis).

Prediction 1. *Distortions induced by inaccurate beliefs cause the slope of the actual investment function to be flatter than the slope of the desired investment function, with the difference in slopes being* $\beta_1(1 - \frac{cov(\tilde{A},A)}{var(A)})$.

To see this, note that this is analogous to the standard measurement error setup,[6] and so the investment relationship can be written as follows:

$$\begin{aligned} s &= \beta_0 + \beta_1 \tilde{A} \\ &= \beta_0 + \beta_1 A + \beta_1(\tilde{A} - A) \\ &= \beta_0 + \beta_1 A + \nu \end{aligned}$$

where $\nu = \beta_1(\tilde{A} - A)$. In general, $cov(A, \nu) \neq 0$, and so the measurement error will cause the slope of the actual investment function to differ from β_1: the slope will be $\beta_1 \frac{cov(\tilde{A},A)}{var(A)}$. Although the magnitude of the slope depends on the exact relationship between $\tilde{A}$ and A, in most beliefs data (as is the case in the data I will present in this paper), measurement error is negatively correlated with A, and so $\frac{cov(\tilde{A},A)}{var(A)} < 1$, resulting in an attenuated slope

[6] One difference from the standard measurement error setup is that, there, only the econometrician has the mis-measured regressor and so the bias is only in the estimated regression line, whereas here, the actor herself is using the mis-measured regressor, and so the bias affects the true investment function.

($|\beta_1 \frac{cov(\tilde{A},A)}{var(A)}| < |\beta_1|$).[7] The intuition behind the attenuated slope is that investments do not respond as much to true achievement as they would optimally.

Estimation

It is difficult to empirically estimate the difference between the slopes of the actual and desired investment functions because neither regression line will in general be causal. Assume parents invest according to the model above plus a noise term due to heterogeneous tastes (ε_i) and consider comparing the slopes estimated from the following two regressions:

$$s_i = b_0^P + b_1^P \tilde{A}_i + \varepsilon_i \tag{1}$$

$$s_i = b_0^A + b_1^A A_i + u_i \tag{2}$$

where (from above) we can see that $u_i = \nu_i + \varepsilon_i$ The slope estimated from equation (1) will be $\beta_1 + \frac{cov(\tilde{A},\varepsilon)}{var(\tilde{A})}$ while the slope from equation (2) will be $\beta_1 + \frac{cov(A,\nu+\varepsilon)}{var(A)} = \beta_1 \frac{cov(\tilde{A},A)}{var(A)} + \frac{cov(A,\varepsilon)}{var(A)}$. Thus, the difference in slopes will be $\beta_1 \left(1 - \frac{cov(\tilde{A},A)}{var(A)}\right) + \left(\frac{cov(\tilde{A},\varepsilon)}{var(\tilde{A})} - \frac{cov(A,\varepsilon)}{var(A)}\right)$ and so will only give us an unbiased estimate of distortions due to inaccurate beliefs (i.e., $\beta_1 \left(1 - \frac{cov(\tilde{A},A)}{var(A)}\right)$) if the second term $\left(\frac{cov(\tilde{A},\varepsilon)}{var(\tilde{A})} - \frac{cov(A,\varepsilon)}{var(A)}\right)$ is equal to 0, that is, if unobserved determinants of investments are either uncorrelated or identically correlated with both A and $\tilde{A}$. To give a concrete example of the potential bias: parents who value education more could have both higher education spending and more overconfidence about their children's achievement, which would bias upwards the estimated slope of the desired investment function relative to the actual, and overstate the distortion.

However, now consider an intervention that changes parents' beliefs from $\tilde{A}$ to A, and thus shifts them from investing based on perceptions, $s^*(\tilde{A}) = \beta_0 + \beta_1 \tilde{A}$, to investing based on the truth, $s^*(A) = \beta_0 + \beta_1 A$. Estimating equation (2) with parents who have received the intervention (treatment parents) will now identify $\beta_1 + \frac{cov(A,\varepsilon)}{var(A)}$, whereas estimating it with parents who have not received the intervention (control parents) will identify $\beta_1 \frac{cov(\tilde{A},A)}{var(A)} + \frac{cov(A,\varepsilon)}{var(A)}$. Since the omitted variable terms are now identical between the two estimates, comparing the slope between treatment and control groups will allow us to see whether inaccurate beliefs

[7] Intuitively, the negative correlation results from the fact that beliefs are bounded, and so measurement error can only be negative at the top of the distribution and positive at the bottom. In the sample, $\frac{cov(\tilde{A},A)}{var(A)}$ is 0.3. An example beliefs generation process that would produce negative correlation is Bayesian updating where the prior is the true distribution and the parent receives signals of true ability. The negative correlation would not hold, however, with (1) simple over (or under) confidence where beliefs are just shifted from the truth by a constant, and, (2) classical measurement error, in which $\tilde{A} = A + \varepsilon$ with ε mean-0 white noise. Note that, under these processes, beliefs about percentiles or scores out of 100 would not be contained within the range from 0-100. Misallocations would still be problematic, but we would have to use other tests to detect them.

distort parents' investments at baseline; that is, to estimate $\beta_1 \left(1 - \frac{cov(\tilde{A},A)}{var(A)}\right)$.[8] This leads to the following testable prediction:

Prediction 2. *If inaccurate beliefs distort parents' baseline investments, then the gradient of the actual investment function will be steeper in the treatment group than in the control group, with the difference in slopes equal to* $\beta_1(1 - \frac{cov(\tilde{A},A)}{var(A)})$.

The intuition is that information allows parents to correct their baseline mistakes and invest along the desired investment function.

This paper investigates the hypothesis that heterogeneity in belief accuracy by parent SES causes low-SES parents to earn low returns. Under this hypothesis, the actual investment function would have a more attenuated slope for low-SES parents than for high-SES parents in the control group (i.e., at baseline), but a similar slope in the treatment group, yielding the following testable prediction:

Prediction 3. *Information will increase the slope of the actual investment function more for groups with less accurate baseline beliefs.*

Note that this assumes that the desired investment function (β) does not vary by group.

This approach is focused on heterogeneity in treatment effects (i.e., changes to the slope of the investment functions), not average treatment effects (ATE's), since ATE's will understate the level of distortions. For example, with unbiased inaccurate beliefs, assuming investments are linear in A, there would be no ATE, but returns would still be distorted. However, if parents are overconfident, information could also have a non-zero ATE. (Note that I use the term "overconfidence" to represent upwardly-biased biased beliefs, as opposed to precisely-biased beliefs.)

Extending the model: Uncertainty

The above model assumes that the slope of the desired investment function does not change when parents receive information. However, the slope could change if we allow for uncertainty in beliefs. If control group beliefs are more uncertain than treatment group beliefs, then the slope of the desired investment function in the control group (which I will denote by β_1^P) would likely be smaller in magnitude than β_1. As a result, comparing the slope of the actual investment line between treatment and control groups would now estimate $\beta_1 - \beta_1^P \frac{cov(\tilde{A},A)}{var(A)}$, which is also a measure of distortions. (Note that $\tilde{A}$ now represents the mean

[8] Note that this assumes that parents fully update their beliefs in response to the intervention. If they only partially update their beliefs, then the metric would be weighted downwards by the updating parameter (i.e., if updated beliefs were a weighted combination of A and $\tilde{A}$ with γ the weight on A, then the difference in slopes would uncover $\gamma\beta_1(1 - \frac{cov(\tilde{A},A)}{var(A)})$).

of baseline beliefs.) To assess whether $\beta_1^P = \beta_1$, one can test whether the heterogeneity in the treatment effect by achievement is equal and opposite to the heterogeneity by baseline beliefs.[9] Alternatively, one can test for treatment effects among people whose baseline belief distribution has an accurate mean.

3 Context and experimental design

3.1 Context

Education system

Primary school in Malawi covers grades 1-8. Primary school has technically been free since 1994, but it does involve expenditures. Parents in the study sample spent an average of 1750 Malawi Kwacha (MWK), roughly 10.6 USD or 1% of annual household income, annually per child. The main expenditures are uniforms (580 MWK or 3.51 USD per child per year), informal but required school fees (380 MWK or 2.30 USD per child per year), and supplemental investments such as school supplies, tutoring, and books (790 MWK or 4.79 USD per child per year). Eighty-nine percent of parents have non-zero supplemental investments.

Dropouts are common in primary school: the nationwide primary school completion rate was 35% in 2007 (World Bank, 2010). In a recent World Bank report, the primary reason cited by pupils for dropping out was "lack of interest," cited by 48% of dropouts. Lack of interest may partially reflect poor performance: 40% of the parents in this study's sample with a child who dropped out during primary school report that, when their children dropped out, they no longer liked school because they were performing badly. Poor performance early-on may thus be a barrier to long-run attainment.

Secondary school, covering grades 9-12, is not free in Malawi, and is significantly more expensive than primary school. Annual secondary school fees for government schools range from 5,000 - 10,000 MWK per year (30 - 60 USD, over 4 times the median primary-school expenditures in the study sample). Parents must also purchase uniforms and supplementary supplies. Anecdotally, many children do not attend secondary schooling as a result of the high fees. Secondary schooling is not open to all students, with admissions strictly governed by an achievement test administered at the end of primary school.

[9] To see this, note that parent i with baseline perceptions $\tilde{A}_i$ and a child with true performance A_i will have a baseline investment of $s^P(\tilde{A}_i) = \beta_0^P + \beta_1^P \tilde{A}_i + \varepsilon_i$. After receiving information, her investments will become $s(A_i) = \beta_0 + \beta_1 A_i + \varepsilon_i$. As a result, the treatment effect as a function of A and $\tilde{A}$ can be expressed as $\tau(A_i, \tilde{A}_i) = s(A_i) - s^P(\tilde{A}_i) = (\beta_0 - \beta_0^P) + \beta_1 A_i - \beta_1^P \tilde{A}_i$. Thus, the heterogeneity in the treatment effect by A will identify β_1 and heterogeneity by $\tilde{A}$ will identify $-\beta_1^P$. (This assumes both β_1 and β_1^P are constant across the sample.) Note that this test works even with incomplete updating, in which case the heterogeneity by A identifies $\beta_1\lambda$ (where λ is the updating parameter) and the heterogeneity by $\tilde{A}$ identifies $-\beta_1\lambda + (\beta_1 - \beta_1^P)$.

School report cards

Schools are supposed to send report cards to parents each term with children's achievement test results. The report card formats vary by school, but all are supposed to have children's absolute achievement test scores and the absolute grades those scores correspond to on the standard Malawian Ministry of Education grading scale of 1-4. (Online Appendix C contains a sample report card used by schools in the study sample.) However, according to baseline survey data, roughly 65% of parents report not knowing their child's performance on his or her last report card, with the main reasons being that the parents (a) did not receive the report card, or (b) were unable to read the report card or understand what the numbers on it meant. Since the reports are supposed to be hand-delivered by students, the former could result from children losing the reports or choosing not to deliver them, as parents of students who performed badly on the last report card were much less likely to receive it.

3.2 Experimental design

The basic idea of the experiment is to gauge parents' beliefs about their children's achievement, deliver true achievement information to randomly selected parents, and then measure the effects on their investments in education. The best individual-level characteristic to focus on for the experiment would be the one that most closely proxies for cross-person variation in returns; in this context, according to qualitative interviews, parents think achievement is the most relevant measure for optimizing their investments and the closest proxy for returns.[10]

Sample selection

The study worked with 39 schools in two districts in central Malawi (the Machinga and Balaka districts).[11] The study team first conducted a sibling census, mapping out the sibling structures for all students enrolled in grades 2-6. Multiple-sibling households were used as the sampling frame because, as described below, I wanted to use inter-sibling tradeoffs to understand parents' long-run investment allocations.[12] The team also gathered term 1 and term 2 achievement test data for the 2011-2012 school year.

[10] Achievement also has the advantage that it determines progression through school and selection into secondary school, thereby almost surely affecting the returns to investment. Some people might think that "innate" ability would have better proxied for returns, but, as has been extensively documented in the literature, it is difficult to measure "innate" ability, so any measure would represent some combination of innate ability and past inputs (e.g., Paxson and Schady (2005)).

[11] Schools were selected randomly from the universe of primary schools, oversampling schools with high and low expected levels of parent education to try to increase the heterogeneity in parent education within the sample.

[12] Using multiple-child households does not have much cost in terms of external validity since fewer than 3% of households in Malawi with children have only one child. The greater potential concern is that households with tighter birth spacing would be over-represented in the sample, but, reassuringly birth spacing is uncorrelated with belief accuracy in the sample, both within and across parent SES categories.

Based on the achievement and sibling data, a sample of 3,464 households with at least two children enrolled in grades 2-6 with achievement test data was drawn. For households that had more than two children that met the sample inclusion criteria, two children were randomly selected for the sample. Thus the sample comprises 6,928 parent-child pairs.

Randomization

I randomly assigned half of the households in the sample to a treatment group that received information about their children's achievement, and half to a control group, which did not receive information.[13] Within the treatment group, half of households were assigned to a "detailed skills" treatment group which received more detailed information about their children's performance (described more below).

Eligibility interviews

Sample selection was based on data gathered from students at school as well as school administrative data. Household eligibility (i.e., whether both siblings lived in the household and were still enrolled in school) was then verified by surveyors who conducted an eligibility questionnaire with parents.[14] Among the 3,464 sampled households, 21% of households were found to be ineligible during the parent interviews, leaving a sample of 2,716 eligible households. Of the 2,716 sampled and eligible households, 97% (2,634 households) were found, available, and consented to participate in the baseline survey. Thus the final experimental sample comprises 5,268 parent-child pairs (observations). Both eligibility and baseline survey completion are unrelated to treatment assignment.

Baseline survey visit

All sampled households were visited by a surveyor who, after verifying household eligibility, asked to speak with the parent who is the primary decision-maker about education.[15] Surveyors then gathered baseline data, including baseline beliefs about both sampled children's achievement, and baseline education spending on each child.

During the elicitation of baseline beliefs, surveyors explained the grading scale used by Malawian schools to parents, including a review of a sample report card which had the same format as the report cards later delivered to the treatment group. This held constant across

[13] The randomization was stratified based on proxies for parent education (both school and, within school, principals' estimates of the maternal literacy rate in a family's village), and a measure of student achievement (the between-sibling achievement gap).

[14] Eligibility for the initial sample was based on children's reports, and so most ineligibility results from misreports by the children. The most common source of ineligibility was that the two sampled children were not siblings but just friends.

[15] If that parent was unavailable, the surveyor spoke with the second parent if there was one and he or she was knowledgeable about educational decisions; if not, the surveyor returned later.

the treatment and control groups both knowledge about the Malawian grading scale and whether parents saw a report card of the type used in the intervention. The specifics of the beliefs elicitation are described in Section 3.6.

After the baseline survey was conducted, during the same visit, surveyors conducted the information intervention (treatment group only).

Information intervention

For the information intervention, surveyors walked treatment parents through two report cards describing the achievement of their two children. The order of presentation was randomized. Each report card contained a child's average performance across all tests administered during term 2 of the school year in his class, specifically, the percent score (an absolute measure), the grade that that score corresponded to on the standard Malawian grading scale, and the within-class percentile ranking.[16,17] These statistics were listed for the three subjects Malawian educators deem most important – Math, English, and Chichewa, the local language – as well as for overall performance (the average across those three subjects). See Appendix A for a sample report card.

The report card also told parents how many achievement tests were included in the averages displayed. On average, students in the sample took 4.5 tests. The correlation between the first and second tests taken was roughly 0.8 for overall achievement, and 0.6 - 0.7 within subjects.

Surveyors walked treatment parents through every number on their report cards. They had received training on how to explain the information clearly. The report card format was chosen based on a series of focus groups and qualitative interviews; the primary criterion for selection was how easily uneducated and illiterate parents could understand the information. After the intervention, the parents were allowed to keep the report cards for future reference.

[16]The tests included "continuous assessments," which were periodic exams administered during the term, and terminal exams, which were administered at the end of the term. For both types, test questions are chosen by teachers from lists of standardized questions contained in the standardized curriculum books given to all schools by the Malawi Ministry of Education. To create the averages used in the report cards and all of the analyses that follow, I use the Malawian Ministry of Education's grading guidelines to create weighted averages, where the weights are 40%/60% (grades 5-6), 60%/40% (grades 3-4), and 100%/0% (grade 2) for continuous assessments and terminal exams, respectively. If a class only offered continuous assessments (or terminal exams), the score used is 100% continuous assessments (or terminal exams). All continuous assessments were combined into an unweighted average. If a student missed an exam, it was not included in their average: parents were informed of this and informed that it could lead to bias in their child's score if tests varied in difficulty and their child missed a particularly easy or hard exam. This could differ from the method used by teachers, who sometimes will replace a child's score with a 0 if they missed the exam.

[17]To be precise, the intervention used "position ranks" instead of percentiles, where the position rank is 100 minus the percentile, a statistic which is easier for parents in this environment to understand than percentiles given a long history of position rankings in schools. Parents' beliefs were also elicited about these "position ranks." For simplicity, I refer to this relative ranking as a "percentile" for the remainder of the paper, and convert position rankings to percentiles for the analysis.

Within the treatment group, parents randomly assigned to the "detailed skills" treatment group also received an additional report card that discussed children's performance on a series of 6 grade-specific competencies (2 math, 2 English, and 2 Chichewa) chosen by local teachers as important skills for children in that grade (See Online Appendix A for sample). The grades displayed on the report card were assigned by the student's teachers. The point of this intervention was to test whether providing more details could help uneducated parents to become more engaged with their children's education, which I will study in other work. For the analysis in this paper, I ignore this subtreatment and pool the treatment households.

3.3 Baseline characteristics of sample

Table 1 presents baseline sample characteristics and tests for balance across the treatment and control groups. Among respondents, 77% are female, and 92% are the primary decision maker about education in the household. Almost half the respondents are farmers. The level of education among parents is very low, with parents having an average of 4.7 years of education. Households are large, with an average of 5 children. The sampled children were in grades 2-6 and were 11.6 years old on average, primarily in the age range from 8 to 16 years old (the 5th and 95th percentiles), with 51% of the children female.

To test balance, I regress each dependent variable on a dummy for being in the treatment group. The differences between the treatment and control groups are never large, and none of the 39 variables tested are statistically significant at the 5% level except for students' baseline math achievement, which is a little lower in the treatment group. This is unlikely to confound the treatment effect estimates since the results mainly look at heterogeneity in treatment effects by child achievement. However, to ensure this is not affecting the estimates, I control for a student achievement measure in all of the regression results. Reassuringly, the accuracy of parents' perceptions (i.e., the absolute value of the gap between true and believed achievement) is balanced across treatment and control groups.

3.4 Data

I use data from two main sources: survey data collected from parents, and administrative data gathered from schools.

Baseline survey data

The baseline survey was conducted from April to June of the 2011-2012 school year. (Term 2 of the school year ran from January - March.) The survey included modules on demographics, income, baseline spending on education, beliefs about the returns to education, and beliefs about children's achievement, described more in Section 3.6.

Immediate investments and updated beliefs

After the baseline survey was conducted and treatment parents received their report cards, surveyors measured three "lab-in-the-field" style investments which I will describe in more detail in Section 3.5: remedial textbooks, level-specific workbooks, and secondary school lottery tickets. Surveyors also measured parents' updated beliefs, described below in Section 3.6. For budgetary reasons, outcomes were measured immediately after the baseline survey and information intervention.

The primary advantage of the "lab-in-the-field" investments is that they were designed to have clear predictions for the perceived "right choice" (i.e., whether the investments are perceived complements or substitutes with achievement). As outlined in Section 2, the clear predictions allow us to easily test for the presence of distortions due to inaccurate beliefs. These investments also allow us to more easily test for heterogeneous responses to information by family background (Prediction 3 of Section 2) because the production function and desired investment functions are likely similar for poor and rich households.

Medium-run investments: School participation and end-of-year grades

In July, 2012, at the end of the third term of the school year, we collected data from teachers' attendance books of students' attendance in the weeks following the baseline survey and end-of-year grades. Due to technical issues with the data collection, we were only able to collect this data for a subset of schools.[18]

Longer-run investments: 1-year household follow-up survey

An endline survey measuring outcomes such as expenditures, dropouts, and other educational investments was conducted approximately one year after the baseline data was gathered (June-July 2013). For budgetary reasons, only a subset of the sample was surveyed.

Unlike the lab-in-the-field investments, for the medium-run and longer-run investments, we do not have as clear *ex ante* predictions about the perceived complementarities. However, these investments allow us to test whether the effects seen with the lab-in-the-field investments are relevant for parents' broader investment decisions. In addition, the experiment also allows us to infer the (perceived) complementarity based on parents' responses.

[18]Specifically, for some schools, the grade and/or attendance were not matchable to the rest of our data. Other schools did not collect the attendance data (grade data are always recorded systematically). The 70% of schools for which we have grade data and 30% of schools for which we have attendance data do not differ on other observable characteristics from the full sample.

3.5 Measurement of immediate investments

Remedial textbooks and level-specific workbooks

First, the survey team measured parents' willingness to pay (WTP) for subject-specific textbooks in Math and English. The textbooks are remedial (i.e., perceived substitutes with achievement),[19] and so the prediction to test is that textbook WTP decreases as parents find out their children are doing relatively better in a given subject. WTP was evaluated using a Becker-DeGroot-Marschak (BDM) methodology, which gives respondents an incentive to report truthfully. See Online Appendix B for a sample price list and description of how the BDM mechanism was implemented.

Parents also make many non-monetary investments that may depend on children's achievement (e.g., asking a sibling to help a child with his homework). For credit-constrained parents, these non-monetary margins might be the primary adjustment margin. To capture this type of investment, surveyors gave parents the choice between receiving 3 different subject-specific workbooks that were targeted for different achievement levels: remedial, average, or advanced. Each parent was given 4 workbooks (one in math and one in English for each of their two sampled children), and, for each workbook, chose which level they wanted to receive. I examine how parents' choices correspond to their children's achievement levels.

Note that, besides having clear predictions for the perceived "right choice" (i.e., production function), an additional advantage of the textbooks and workbooks is that there are clear predictions for the actual right choice. For example, the advanced workbook was designed specifically to be better for higher-achieving children in the study sample. This enables me to make a stronger argument that parents' misallocations lower actual returns, not just perceived returns. For the lottery, described next, I must rely on estimates from the literature on complementarities in other contexts.

Secondary school lottery

Secondary schooling is one of the first high-cost investments that parents make, so we measured a short-run, real-stakes proxy for secondary schooling investments. Specifically, we conducted a lottery to pay for four years of government secondary school fees for one child in every 100 households in the sample; four years of fees were worth roughly 120 - 240 USD at the time of the experiment. Parents were given nine tickets for the lottery and chose how they would allocate the tickets across their two sampled children.[20] There are two primary

[19] Before the study began, we surveyed teachers at schools in the sample and 100% of teachers surveyed thought that the textbooks were more useful in subjects in which children have lower achievement than in subjects in which children have higher achievement.

[20] There are two reasons that I used multiple tickets instead of one ticket: First, in a setting with inequality aversion, it increases my power to detect small shifts that would be inframarginal if there were just one ticket,

ways that student achievement might affect the expected return of a lottery ticket. First, the labor market returns to completing secondary school might be higher for higher-achieving students. Most parents in this setting believe that is the case (i.e., that achievement and secondary schooling are complements). Second, the probability of admission to secondary school increases with achievement: at the end of primary school, students take a standardized achievement test, which is the sole determinant of secondary school admissions.[21] Therefore, both the expected value of the fees paid by the scholarship and the probability of receiving the secondary school earnings return should increase with achievement.

3.6 Measurement of baseline and updated beliefs

To measure parents' baseline beliefs, surveyors asked parents about the same achievement metric that was delivered to treatment parents in the intervention – average achievement on all school exams from term 2 of the schoolyear.[22] I used the same measure that was later delivered to treatment parents, instead of asking more broadly about academic abilities, for two reasons: first, so that any divergence between elicited beliefs and true achievement would represent inaccuracies in beliefs instead of differences between measures, and, second, because inaccurate beliefs about achievement are instructive since parents say it is their primary metric for optimizing. The beliefs were elicited for overall, math, English, and Chichewa scores for both sampled children, and then repeated for percentiles. The uncertainty of parents' beliefs was also measured by asking parents to distribute tokens across bins representing different achievement ranges (e.g., 0-20, 20-40). Beliefs were also elicited about all specific competencies covered on the "detailed skills" report card.[23]

To assess how much parents updated their beliefs, at the very end of the survey (after the treatment group had received information), surveyors asked respondents how well they thought that their children would perform on an achievement exam if they took it that same day. Note that, unlike baseline beliefs, this is *not* the same measure that the treatment group received information about.[24]

and second, it allows me to also use this lottery to study inequality aversion (the subject of a different paper).

[21] The lowest performers fail the exam and are ineligible for any secondary school (25% of test-takers in 2010); the next tier pass but are not admitted to government school; the highest performers are admitted to government secondary schools. The lottery guidelines stated that it would pay for the full fees if the child were selected into government school, and 0 otherwise.

[22] Parents were given a visual aid with a line showing all scores (at 5 point increments) from 0 - 100. They were asked to point to the score that they thought was closest to their child's average score on all exams offered during term 2 of that year; beliefs were thus elicited in 5-point increments.

[23] For each skill (e.g., adding 2-digit with 2-digit numbers), parents were asked whether their child could do the skill, with the possible responses being 1=Yes, 2=A Little, 3=No, 4=Don't know, 5=Can't understand skill.

[24] Although re-asking about term 2 exams would have given a more direct measure of whether parents understood the information that they had been given, that would not have shown us how much parents

4 Baseline results: Parents have inaccurate perceptions that affect their investments

4.1 Gap between perceived and true achievement

I first examine whether parents have inaccurate perceptions by comparing parents' baseline beliefs about their children's achievement on the exams administered at their schools the previous term with their children's true achievement on the same exams.[25] The left figure in Figure 3 shows kernel density plots of parents' beliefs about their children's overall test scores (the solid line) and children's true test scores (the dashed line). Scores are absolute scores expressed on a scale from 1 to 100.[26] Parents are overconfident, but, beyond simple overconfidence, beliefs are more broadly inaccurate: the right figure plots a kernel density plot of each individual parent's beliefs relative to their child's true achievement. If parents were simply overconfident by, say, 5-10 points, the plot would have all of its density between 5 and 10; rather, the density is spread widely. Twenty one percent of parents are underconfident.

The magnitude of the gap between perceived and true achievement is large: 20 points, or 1.2 standard deviations of the achievement distribution on average. The gaps for subject-level beliefs are slightly larger (Table 1). This corresponds to correlations of 0.3 between believed and true achievement for overall achievement, and 0.2-0.3 for subject-level. Recall that the correlation between the tests the student took during the term was much higher (0.8 for overall achievement, and 0.6 - 0.7 within subjects).[27] This suggests that the disconnect

updated their underlying beliefs about their children's academic abilities, since parents should have fully updated about that measure even if they did not update their underlying beliefs about their children's full set of abilities – beliefs which had been developed over multiple years and likely have tighter priors.

[25] Elicited beliefs are binned at 5-point increments; results are robust to binning true achievement scores at 5-point increments as well.

[26] I focus on absolute performance information for two reasons. First, parents appeared to respond more to absolute than to relative performance (e.g., if one simultaneously analyzes how a parent responded to the shock to their absolute and to their relative beliefs, parents on average responded more to the shock to their absolute beliefs). Second, there was an implementation issue with the relative achievement information delivered to the first 595 treatment households. All of the absolute performance information they received was correct, but they received two pieces of incorrect relative performance information: for one child, in the space for true overall relative performance, their Chichewa relative performance was listed (which has a correlation of 0.83 with the true overall), and for the other child, in the space for math relative performance, their English relative performance was listed (correlation of 0.55 with the true math). The results are robust to dropping the treatment households (and corresponding controls) that received the incorrect relative performance information, and to using either absolute or relative performance for the analyses. The households given incorrect information were later revisited to give correct information.

[27] I can also benchmark it against the correlation between current-term achievement and previous-term achievement, which are 0.6 for overall and 0.5 for subject-level. The data for term 1 and term 2 are only matchable for a subset of the sample.

between beliefs and achievement does not simply reflect noise in achievement. Parents also have inaccurate beliefs about between-subject (math relative to English) achievement, and inter-sibling relative achievement (child 1 relative to child 2), with beliefs about the inter-sibling achievement gap deviating from the true gap by an average of 19 score points (1.1 std. dev.), and 31% of parents wrong about which of their children has higher scores. Misunderstanding the difficulty of the grading scale does not seem to drive inaccurate beliefs, as Appendix Figure 1 shows similar patterns using within-class percentiles.

Since achievement tests determine progression through school, and since inaccurate beliefs about achievement correlate with inaccurate beliefs about child educational competencies (e.g., whether the child can do 2-digit addition), these inaccurate beliefs are likely relevant for a broad range of educational investments, as shown in the subsequent sections.

4.2 Desired vs. actual investment functions at baseline

Motivated by Theoretical Prediction 1, Figure 4 presents suggestive evidence on whether inaccurate beliefs cause distortions by comparing the slope of the desired investment function (investments plotted against believed performance – the dashed lines) with the slope of the actual investment functions (investments plotted against true performance – the solid lines). Note that the y-axes for both lines represent investments, but the x-axes differ. Both lines are locally linear regression lines. Recall that the prediction is that, if there are distortions, the actual investment function will be flatter than the desired investment function. The data are from the control group only to evaluate how investments are distributed in the absence of information. I begin by interpreting the desired investment functions and then move on to the actual.

Desired investments: Workbooks (complements) and textbooks (substitutes)

Panel (a) shows the graphs for math and English workbook choices: here, the y-axis represents the parents' choice between beginner/average/advanced workbooks, with the 3 choices parametrized as -1/0/1 for simplicity (so parents who choose a beginner workbook are coded with -1, average with 0, etc.).[28] As expected since workbook difficulty is a perceived complement with achievement, the dashed line slopes upwards, showing that parents choose more difficult workbooks when they believe their children are performing better.

Panel (b) shows the relationship for WTP for remedial textbooks. I analyze between-subject WTP (i.e., math WTP - English WTP, equivalent to running with child fixed effects (FE's)) because it has more clear predictions for behavior than within-subject WTP:

[28] The relationship is robust to other parametrizations; e.g., indicators for choosing the beginner workbook, an indicator for choosing the advanced workbook. Recall that all workbooks are free; this choice shows us how parents tailor their non-monetary investments.

between-subject WTP holds constant parents' other investments in education, allowing us to isolate whether they spend relatively more in the subject in which they think their child is underperforming.[29] The y-axis shows the log of WTP for the remedial math textbook minus the log of WTP for the remedial English textbook.[30] The x-axis shows the parents' beliefs about the child's performance in math relative to English. Because the textbooks are remedial (substitute with achievement), the prediction is that parents are willing to pay more for the book in the subject their child is more behind in, which is what we see since the line slopes steeply downwards.

The relationships between perceived achievement and investments are highly statistically significant: linear regression coefficients yield t-stats of 33, 43, and 15 for math workbooks, English workbooks, and textbooks, respectively.

Desired investments: Secondary school lottery

Panel (c) of Figure 4 shows how parents' ticket allocations in the control group depend on children's achievement. For the dashed line, the dependent variable is the number of secondary school lottery tickets given to the child the parent perceives is higher-achieving minus the number given to the child she perceives is lower-achieving, and the independent variable is the perceived performance gap between the perceived-higher-achieving child and perceived-lower-achieving child.

For interpreting the results, it is useful to note that over 75% of parents split their nine lottery tickets as evenly as possible, consistent with a high degree of inequality aversion (see Appendix Figure 2). Thus, in most cases, I am analyzing which child parents give their ninth ticket to, which is analogous to parents' real-world decisions when they have to choose between their children (e.g., if they can only afford to send one child to secondary school).

The line slopes upwards, and remains above 0 throughout: across the treatment group, parents are giving more tickets to the child they think is higher-performing.[31] The slope

[29] To see that the within-subject predictions are less clear, consider a parent who has received negative information about their child's math achievement. Because the math textbook is remedial, holding all else constant, the parent's WTP for the math textbook should increase. However, all else is not held constant: the negative shock to math performance is correlated with a negative shock to overall performance, which means that, say, the parent might increase her estimate of the chances that her child will drop out of school in the next year, thereby decreasing the value of the textbook. The net prediction is ambiguous. In contrast, consider a parent who received negative information about how well their child was performing in math relative to English. In this case, comparing their math WTP with their English WTP would hold constant the parent's estimated chances of child dropout, and give the unambiguous prediction that math WTP should fall relative to English WTP.

[30] Logs are used to improve precision, but results are robust to using levels. Only 6% of WTP observations have values of 0; for these, I replace with the log of 10% of the lowest value of the price list, but, since there are so few 0's, the results are nearly the same if I drop the 0's from the regressions.

[31] Online Appendix Table A.1 provides evidence that this is not just because beliefs are correlated with other factors by showing that the predictive power of perceived achievement is robust to controlling for child

is steepest near 0, which could reflect that parents' decisions are constrained by inequality aversion.

Desired vs. actual investment functions

I now compare the slope of the actual investment functions (solid lines) with the slope of the desired investment functions. For the workbooks and textbooks, the solid lines have the exact same y-axis as the dashed lines, but a different x-axis. For the lottery, the dependent variable for the solid line is now the number of lottery tickets given by the parent to the *true* higher-achiever relative to the true lower-achiever, and the independent variable is the true achievement gap between the higher-achieving and lower-achieving child.

The results show that the gradient of investments on true achievement is flatter than the gradient on beliefs. For the textbooks and workbooks, although the solid lines all have non-zero slopes (statistically significant at the 5% level in linear regressions), the coefficient magnitudes are much smaller than for the dashed lines, only 20-33% of the magnitude. For the lottery, the solid line is flatter than the dotted line, and everywhere below it, with the difference between lines statistically significant everywhere except for near zero and large positive values on the x-axis, where density is low. Parents appear to be trying to give more tickets to their higher-performing child, but are prevented from doing so since they do not always know who their higher-performing child is.

The finding that investments have a steep gradient with beliefs but a much flatter gradient with the truth suggests that parents try to tailor their investments to their children's achievement, but are prevented from doing so by inaccurate beliefs. However, as outlined in the theoretical predictions, this evidence is suggestive, not causal: both beliefs and achievement could be correlated with other factors that determine parents' investment decisions. For example, parents who have a preference for a given subject might be overconfident about their children's achievement in that subject and might also invest more (difficulty or money) in that subject. This could produce results like we see for the complements (workbooks) but not the substitute (textbooks) investments. Another variant is that factors other than children's achievement (e.g., affection for the child, the child's work ethic) underlie parents' decisions and are more highly correlated with beliefs than true achievement. Although it is reassuring that the slopes conform to the causal predictions laid out in Section 2, stronger evidence is needed. The information experiment can help establish whether the difference is causal: under the omitted variable stories, investments would not respond to information because control parents would not be making "mistakes." In contrast, if there are misallocations due to inaccurate beliefs, then investments would respond.

gender or age.

5 Impact of information

In this section, I test whether information increases the gradient of the actual investment function for workbooks, textbooks, and the lottery.

5.1 Effect of information on beliefs

The information experiment only has power to affect investments if it changes parents' beliefs. Figure 5 compares the absolute value of the difference between each child's term 2 achievement and her parent's beliefs between the treatment group and control group before vs. after the intervention. The dark grey bars use the beliefs elicited at the beginning of the baseline survey (pre-intervention) about children's term 2 achievement; the light grey bars use the beliefs elicited post intervention about how well the child would do on an achievement test if he/she took it that day. (Note that, for reasons described in Section 3.6, the baseline and post-intervention beliefs questions asked about different measures.) In the control group, the baseline and post-intervention beliefs are similarly far from children's true term 2 achievement, but, in the treatment group, the post-intervention beliefs are 7.6 points closer, which is significantly different from both the control group's post-intervention beliefs and the treatment group's baseline beliefs (p-value $< .01$). Note that respondents do not fully update their beliefs, potentially because they are Bayesian updaters who have prior information about their children's underlying abilities.

5.2 Treatment effects on the slope of the actual investment function

Graphical evidence

I now use the information experiment to test whether the differences in Figure 4 between the investment gradient on the truth and beliefs represent the causal impact of inaccurate beliefs. Figure 6 shows local linear regression plots of parents' actual investment functions: investments on the y-axis against true achievement on the x-axis. The solid line represents the control group (so is the same as the solid line from Figure 4); the dashed line represents the treatment group.

Per Theoretical Prediction 2 from Section 2, if the differences in Figure 4 between the control group's desired and actual investment functions represent the causal impact of inaccurate beliefs, then the treatment group's actual investment functions should be much steeper than the control group's, with a slope more similar to that of the control group's desired investment function (dashed line in Figure 4). In contrast, if the differences in Figure 4 were a non-causal result of correlations with omitted factors differing between truth and

beliefs, then the treatment line should look the same as the control line.

Consistent with a causal interpretation, Figure 6 shows that the information treatment increases the slope of the treatment groups' investment function.

Regression evidence: Textbooks and workbooks

I now perform a formal test of whether information changes the slope of the investment functions (Prediction 2) by running the following regression:

$$s_{ij} = c_0 + c_1 A_{ij} \times Treat_i + c_2 A_{ij} + c_3 Treat_i + c_4' X_{ij} + \varepsilon_{ij} \tag{3}$$

where i indexes households, j indexes siblings, s is the investment, A is the relevant achievement metric (e.g., math for math workbooks, math - English achievement for between-subject textbook WTP), $Treat_i$ is an indicator for being assigned to the treatment group, and X_{ij} is a vector of control variables.[32] Since each household has multiple observations (one for each sibling $j \in \{1, 2\}$), standard errors are clustered at the household level.

The prediction is that the information treatment makes the slope steeper, so that $c_1 > 0$ for complements (the workbooks), and that $c_1 < 0$ for substitutes (textbook WTP). (Note that c_1 will provide a measure of the $\beta_1(1 - \frac{cov(\tilde{A},A)}{var(\tilde{A})})$ metric from Section 2 which allows us to assess how much inaccurate beliefs distort investments.)

Table 2 presents the regression results. Panels A and B use math and English workbook choice as the dependent variables, and Panel C uses the log of WTP for the math remedial textbook minus the log of WTP for the English remedial textbook. Column (1) shows the base specification: consistent with the graphical evidence and the predictions of the model, c_1 is positive for the workbooks and negative for the textbooks. All 3 coefficients are highly statistically significant (p< .001). Moreover, the magnitude of the effects is large: comparing the coefficient on *Score* (slope in the control group) with the sum of the coefficients on *Score* and $Treat \times Score$ (slope in the treatment group) shows that parents' investments were 3, 2.7, and 5.3 times more responsive to a given change in achievement in the treatment group relative to the control group for math workbooks, English workbooks, and textbook WTP, respectively. These investments were chosen specifically to allow us to see whether inaccurate beliefs cause misallocations; the levels of the investments – and thus the ATE's– are not particularly interesting in and of themselves. However, for completeness, col. (6)

[32]Results are robust to excluding the controls. Control variables include school fixed effects (FE's), the between-child score gap, and parents' education level. Note that this includes all variables underlying the stratification but not the stratum fixed effects themselves; I do not control for stratum FE's because some of the stratum are very small, and so 20% of observations would be lost because there is no variation in treat within their stratum. This decision was specified before the experiment was conducted. The results are, however, robust to controlling for stratum FE.

shows the ATE's.[33]

Regression evidence: Secondary school lottery

Because the number of lottery tickets was constrained at the household level, the regression analysis includes a household fixed effect, and assesses the number of tickets given to one child relative to his sibling. I thus run the following regressions:

$$Tix_{ij} = c_0 + c_1 Treat_i \times 1\{HigherPerformingSib\}_{ij} + c_2 1\{HigherPerformingSib\}_{ij} + \tau_i + \varepsilon_{ij} \tag{4}$$

$$\begin{aligned} Tix_{ij} = & c_0 + c_1 Treat_i \times 1\{HigherPerformingSib\}_{ij} + c_2 1\{HigherPerformingSib\}_{ij} + \\ & c_3 Treat_i \times A_{ij} + c_4 A_{ij} + \tau_i + \varepsilon_{ij} \end{aligned} \tag{5}$$

where Tix_{ij} represents the number of tickets given to sibling j in household i, $1\{HigherPerformingSib\}_{ij}$ is an indicator that sibling j is the higher-achieving sibling in his or her house, A_{ij} is child j's achievement, and τ_i is a household fixed effect. Equation (4) allows us to see whether the treatment causes parents to shift tickets towards their higher-achieving child: the prediction is that $c_1 > 0$. Equation (5) tests for whether the size of the effect depends on the performance gap between the children: the prediction is that $c_1 > 0$ and/or $c_3 > 0$, depending on whether parents primarily care about rank order or the performance gap.

Panel D of Table 2 shows the regression results. Column (1) shows the regression of equation (4). Information cases parents to allocate an average of 0.98 more tickets to their higher-scoring sibling relative to their lower-scoring sibling (t-stat=7.5), a large effect on a base of 0.53. Column (2) tests for whether the information also changes the slope of the line (equation (5)), and finds that there is no significant heterogeneity: parents seem to primarily use the rank order information, which is a logical way to make the decision if, because of inequality aversion, parents are splitting their tickets 5-4 and just choosing which child to give 5 tickets to.

Regression evidence: Robustness

One potential concern with the previous analyses is that performance is obviously not randomly assigned. Thus, if there is heterogeneity in the effect of information based on some other factor correlated with performance, then it could also cause a change in the slope.

[33] ATE's can result from overconfidence (negative/positive ATE's for complements/substitutes) and uncertainty (positive ATE's for both). Here uncertainty leading to underinvestment is likely minor since parents were forced to make choices, and so, unsurprisingly, the ATE's reflect overconfidence (parents were overconfident about math, English, and math relative to English at baseline).

It is reassuring that the direction of the effects fits exactly the theoretical predictions from Section 2, combined with the baseline investments analysis from Section 4. In addition, columns (2) through (5) provide evidence that omitted variables do not drive the result by showing that the results are robust to controlling for household fixed effects (panels A-C only – all specifications already include household fixed effects in panel D) and to controlling for interactions of individual-level control variables with treatment.[34]

5.3 Additional treatment effects analysis

This section contains additional analysis to try to gain a fuller sense of the causal chain between information, beliefs, and investments, as well as a brief analysis of gender effects for the lottery.

Uncertainty

In addition to changing the mean of beliefs, information may also decrease the uncertainty of beliefs, which could affect investments if the slope of the desired investment function depends on uncertainty ($\beta_1^P \neq \beta_1$). In Appendix Table 3 (col. (4), Panel A), I test for uncertainty effects by examining the treatment effects among parents whose initial beliefs are close to their child's true achievement. The main uncertainty effects are for the lottery: treatment parents who received confirmation that their initial beliefs were essentially correct allocate 0.5 more tickets to the child that is higher performing.[35,36] Note that an alternative explanation is that the treatment increased the salience of achievement.

Asymmetric impacts for positive vs. negative information shocks

Parents' responses to information might depend on whether they receive a positive information shock (i.e., true achievement is higher than their beliefs) or a negative one. Since people like good news, we might expect them to respond more to positive shocks. The following regression estimating heterogeneity in the treatment effect by the information shock

[34]Note that, if all we wanted to do was identify whether information caused investments to be more closely aligned with performance, then random assignment of the treatment is sufficient for identification.

[35]See also Appendix Table 2 which uses the second test described in section 2, with consistent results: for workbooks and textbooks, we cannot reject that heterogeneity in the treatment effect by parents' baseline beliefs is the same as heterogeneity by true achievement, but we can for the lottery.

[36]The presence of uncertainty effects for the lottery may make one wonder if the reason that parents split their tickets so evenly is that they are unsure which child would be the better investment (the alternative hypothesis is that they are averse to investing unequally in their children). I provide some evidence on this by regressing the absolute value of the gap in tickets given to the two children on $Treat_i$. If uncertainty were the primary driver, we would expect parents in the treatment group to split their tickets less equally than parents in the control group. However, I find that the treatment only increased the gap by 0.14 tickets on average (p-value only 0.17); this means 73.5% of treatment parents split their tickets as evenly as possible relative to 75.3% of control groups. Thus, although uncertainty may play some role in high degrees of ticket splitting, it does not seem to be a primary factor: inequality aversion likely is more important.

is informative for this issue:

$$s_{ij} = a_0 + a_1(A_{ij} - \tilde{A}_{ij}) \times Treat_i + a_2 Treat_i + a_3(A_{ij} - \tilde{A}_{ij}) + \varepsilon_{ij} \tag{6}$$

where A_{ij} is child j's achievement and $\tilde{A}_{ij}$ is the parent's baseline beliefs about child j's achievement. The heterogeneity in the treatment effect, a_1, is a scale factor that shows us how much more outcomes change for every additional point of divergence between true and believed achievement. The prediction is thus that a_1 is larger in magnitude for positive shocks (so, more positive for outcomes, like workbooks, that increase with achievement and thus have positive a_1). Appendix Table 1 tests for heterogeneity in a_1 by estimating equation (6) fully interacted with an indicator for whether parents received a positive information shock ($A_{ij} > \tilde{A}_{ij}$). I find that parents respond more to positive shocks.[37]

The lottery and gender

Information could affect the distribution of tickets based on other child characteristics that are correlated with beliefs. One hypothesis is that parents underestimate their daughters' ability and achievement and this is one cause of underinvestment in girls' education. This is not what the results show. If anything, treatment parents allocate fewer tickets to their girls, although the difference is not statistically significant (p-value=0.21).[38]

6 Heterogeneity by parent education

The previous section suggests that inaccurate beliefs distort parents' investments. In this section, I examine heterogeneity by parent education to shed light on the hypothesis that these distortions provide a channel for the perpetuation of inequalities across generations. I first show that less-educated parents have less accurate beliefs. For their larger belief inaccuracies to translate into larger treatment effects on their investments, less-educated parents would need to update their beliefs more in response to the intervention (i.e., to have a larger gap between updated and baseline beliefs), and so I next verify that this is the

[37] Cols. (1), (3), and (5), show the regression results using math workbooks, English workbooks, and post-intervention beliefs as the outcome variables; I omit textbooks and the lottery since the relevant achievement measure for those is relative achievement and so positive shocks do not correspond to "better news." A potential confound is that parents' behavior is bounded; for example, if a parent would have chosen the beginner workbook in the absence of information, then she could only respond to information by choosing a weakly more-advanced workbook (since there are no less-advanced workbooks). To address this, cols. (2), (4), and (6) restrict the sample to parents whose predicted behavior based on baseline beliefs would be predicted to be near the middle of the range of potential outcome variables. The magnitudes of the results are similar and/or stronger, suggesting that this concern does not drive the results.

[38] This could partially reflect the fact that parents in fact overestimate their girls relative to their boys, as girls are performing worse on average in school (roughly 2 points lower achievement) but parents believe their girls are performing almost as well as their boys (beliefs only 0.45 points lower on average).

case. (Note that, for there to be larger effects on their investments, less-educated parents do not need to update their beliefs more conditional on the size of the gap between baseline beliefs and the truth, just absolutely, although that stronger finding holds as well, as shown below.) Finally, I test whether less-educated parents have larger treatment effects on their investments – and thus larger implied baseline distortions – and find evidence that they do. Note that the statistical power of the analysis is limited by the fact that there are not many educated parents in the sample, but the effects are still significant.

6.1 Belief accuracy

Recall that Figure 1 suggests that less-educated parents in Malawi have less accurate beliefs. Table 3 investigates this more rigorously by regressing parents' belief accuracy (the absolute value of the gap between believed and true achievement) on the average years of education among parents in the household.[39] The raw gaps (shown in the odd-numbered columns) show that less-educated parents have less accurate beliefs, although the findings for English are not significant, perhaps because only the most educated parents in the sample understand English. The findings are robust to child and parent controls (such as gender; see the even-numbered columns), and exist within schools for math and Chichewa (Appendix Table 5, odd-numbered columns), suggesting that less-educated parents do not simply attend schools that give worse information. Appendix Table 4 shows that the heterogeneity is not due to the particular measures used, but is robust across different measures of parent education and child achievement. Since the children of less-educated parents also have lower achievement, one potential explanation is that beliefs are less accurate at lower achievement levels. Note that, because achievement is not exogenous, controlling for it in the regression will likely bias downward the estimated effect of parent education relative to the true effect; for example, less-educated parents' less accurate beliefs could have contributed to their children's low achievement. That is, controlling for achievement is "over-controlling" by conditioning on a downstream variable. However, it is interesting that some significant gaps still remain even controlling for achievement (even-numbered columns, App. Table 5).[40]

[39]The results are robust to using the respondent's years of education (see Appendix Table 4), but the average across parents in the household is used for two reasons: First, focus group discussions held before the project began indicate that parents share information. Second, the data (presented in Online Appendix Table A.2) provide suggestive evidence of both information sharing and information dilution between parents: col. (1) and col. (2) show that both parents' education matter; for mothers (the majority of respondents), if anything, the respondent's spouse's education matters more, although one cannot reject equal effects. Col. (3) shows that the respondent's own education matters more for one-parent households, which is consistent with information dilution, although there are obviously many other differences between one- and two-parent households. As a result, col. (4), and the specifications in the main tables, use the average across parents in the household.

[40]Online Appendix Table A.3 provides further suggestive evidence that the heterogeneity in belief accuracy by parent education reflects heterogeneity in their ability to assess a child's performance. The table shows how

Less-educated parents also have more uncertain beliefs (Appendix Table 6, cols. (4) - (6)), but are not significantly more overconfident (cols. (1)-(3)).

6.2 Updating

Table 4 looks at whether less-educated parents change their beliefs more than more-educated parents in response to information. Column (1) regresses the absolute value of the difference between a respondent's baseline beliefs and post-intervention beliefs about overall achievement on a dummy for treatment and its interaction with parent education. (Recall that baseline beliefs were about term 2 achievement whereas the post-intervention beliefs were about a slightly different metric – how the child would perform on an achievement test taken that day – but that the difference can still proxy for the change in beliefs.) Less-educated parents have a significantly larger treatment effect than more-educated parents. In terms of magnitudes, going from no education to completing primary school decreases the treatment effect by roughly 3 score points, or 54% of the control group mean. The main reason that less-educated parents update more is that they have less accurate baseline beliefs, but, consistent with higher uncertainty, they also seem to have a higher Bayesian updating parameter. Specifically, column (2) controls for the interaction between treatment and the size of the gap between initial beliefs and the truth. The heterogeneity remains sizable and significant at the 1% level. Cols. (3) and (4) look at whether the treatment shifted the mean of beliefs more negatively for less-educated parents; the differences are not statistically significant.

6.3 Treatment effects on investments

I now examine whether heterogeneity in belief accuracy and updating translates into heterogeneous treatment effects on investments. The analysis is complicated by the fact that it assumes the desired investment function (β_1 from the theoretical predictions) does not vary by parent background. This is most plausible for the workbooks, for which it is difficult to see why more educated parents should have a different mapping between achievement and a free choice. This is less likely for textbooks and the lottery; for example, for textbooks, credit constraints could change how much parents can spend and thus how much their spending

perception accuracy about overall (cols. (1)-(2)), math (cols. (3)-(4)), English (cols. (5)-(6)), and Chichewa (cols. (7)-(8)) achievement changes as students progress through school. For Math performance (col. (3)), parents' belief accuracy decreases as students age, which probably results from the material getting more difficult, making it harder for parents to judge performance on their own. However, col. (4) shows that the pattern is less pronounced for more-educated parents. This is consistent with a role for parent judgement, and for less-educated parents having a harder time judging their children's performance as the material becomes more difficult. Note that the performance gap in math does not follow a similar pattern and so does not seem to explain the finding: the children of less-educated parents actually catch up to the children of more-educated parents in math as they progress through schools. For English and Chichewa, we do not see the same pattern, as it may be easier for parents to judge their children's language performance as their children improve and can speak and translate.

responds to achievement.

Figure 7 shows the treatment effect graphs for workbooks (panel (a)), textbooks (panel (b)), and the secondary school lottery (panel (c)), estimated separately using respondents without secondary education (left column) or with secondary and above (right column); the graphs look similar with different education measures. Starting with the math workbooks, there appear to be two differences between the graphs. First, the control (solid) line is flatter for the less-educated parents: they are worse at targeting their investments at baseline. Second, the treatment (dashed) line is steeper for the less-educated parents: this is consistent with them updating their beliefs more. The patterns for English workbooks look similar but noisier and less pronounced (consistent with the heterogeneity in belief accuracy being smaller in English).

Table 5 tests formally for these patterns by estimating equation (3) interacted with the average years of education of parents in the household. Columns (1), (4), and (7) show heterogeneity in the treatment effects for workbooks and textbooks, and columns (2), (5), and (8) show heterogeneity in slopes for the treatment group only. I find that, for both math and English, less-educated parents' workbook choices respond more to information than more-educated parents' decisions do. For example, for math, the baseline slope is flatter for parents with less education and the treatment slope steeper (coefficients on $Score \times ParentYrsofEduc$ in col. (1) and col. (2)). As a result, less-educated parents change their investments more in response to information, with the treatment effect on the slope decreasing by roughly 6% (-.12/1.92) for each additional year of education, significant at the 1% level (col. (1), coefficient on $Treat \times Score \times ParentYrsofEduc$).

For both textbooks and the lottery, the point estimates indicate that the investment gradient of less-educated parents changes more than that of more-educated parents (coefficient on $Treat \times Score \times ParentYrsofEduc$), but the differences are not statistically significant. The heterogeneity in between-subject belief accuracy by parent education is limited so likely explains the textbook results. For the lottery, there is significant (but small) heterogeneity in accuracy about the between-sibling achievement gap, so the results may reflect low statistical power, or heterogeneity in the desired investment functions (e.g., because of heterogeneity in inequality aversion).

7 Results: Medium-run and longer-run outcomes

The above results demonstrate that inaccurate beliefs affect parents' investments in education. One open question, however, is whether the effects of inaccurate beliefs are relevant beyond the controlled survey environment. To examine that, I analyze data from the endline survey conducted with parents roughly one year after the baseline survey, as well as data

collected from schools several months after the baseline. The advantage of these data is that they allow us to gauge the persistence and external validity of the earlier results. However, these results are noisier and harder to cleanly interpret than than those presented above since they reflect other factors, including the reaction of the children to the information, the resulting responses of the parents to the children, etc. Note that, because the intervention was relatively small, just a one-time infusion of information, we can think of these effects as lower bounds on the effects of a larger, more sustained intervention.

7.1 Persistence of beliefs

I first check whether the information affects treatment parents' beliefs one year after the intervention. If they forget the information, then the treatment effects would likely not persist over time. In the end line survey, we elicited parents' beliefs about their children's current achievement. Online Appendix Table A.4 verifies that the endline beliefs of treatment parents correspond more tightly with their children's past achievement than the endline beliefs of control parents.[41]

7.2 Information treatment effects

I now analyze the effect of the information treatment on the slope of the actual investment function.[42] Figure 8 shows graphical evidence for the primary longer-run investments measured.[43] Table 6 shows the regression results. Columns (1) and (2) present coefficients on $Treat$ and $Treat \times A$ from estimation of regression 3; for ease of interpretation, columns (3) and (4) present an alternative specification where the outcome variable is regressed on $Treat$ and $Treat \times 1\{AboveMedianA\}$ (where $1\{AboveMedianA\}$ is an indicator for having above-median achievement). All regressions use the child's overall achievement on the term

[41]Specifically, I regress parents' beliefs at endline on children's true past achievement (i.e., what we delivered to the treatment group in the intervention), $Treat$, and an interaction of past achievement with $Treat$. For beliefs about overall, math, and English achievement, the beliefs of parents in the treatment group are more closely aligned to past achievement (the coefficient on $Treat$ is negative and on $Treat \times Score$ is positive), although the relationship is much stronger and only statistically significant for overall achievement, which was likely the most salient measure for parents. For Chichewa, the relationship in fact goes the wrong direction; this could reflect that parents in the treatment group actually thought they helped their children to improve significantly in Chichewa.

[42]Note that I do not first analyze the control group's beliefs data to get predictions for the desired investment function as I did in Section 4. This is for two main reasons: First, these outcomes are more subject to reverse causality and omitted variable bias than the lab-in-the-field investments, and so the estimated desired investment function will be more biased and difficult to interpret (e.g., it is unlikely that choosing an advanced workbook today would cause a parent to believe her child's past achievement is high, but that type of reverse causality is much more likely for, say, expenditures, which are highly correlated over time.) Second, because parents' beliefs evolve over time, parents' baseline beliefs are a noisy proxy for their beliefs one year later.

[43]Due to smaller sample sizes than for the outcomes examined earlier, the lines are generally not statistically significantly different at a given point and so I remove the confidence intervals for ease of interpretation.

2 2011-2012 achievement exams (the same measure used for earlier regressions) as the measure of A. Finally, column (5) tests for an average treatment effect (ATE). Relative to the investments examined earlier, the level of these investments is more interesting (e.g., we care if the intervention changed overall spending levels more than we care if it changed average WTP for math relative to English textbooks). So, while the primary objective is still to test for misallocations (treatment effects on the slope of the investment functions), it is also interesting to test for ATE's.

Starting with dropouts, information changes the slope of the investment function. High-achieving students in the treatment group are less likely to drop out of school, while low-achieving students are more likely, which is what we would predict given that most parents believe that education and achievement are complements. The change in the gradient is significant at the 1% level. In terms of magnitudes, dropout falls by 1.5 percentage points for students who were above-median achievement and increases by 2.2 percentage points for students who were below-median (cols. (3) and (4)). These are large effects relative to the control group mean (2%). However, there is no statistically significant ATE on dropouts in the treatment group. Since parents were overconfident at baseline, the fact that we do not see a corresponding increase in dropouts could reflect uncertainty effects (i.e., decreased uncertainty increasing investments) or asymmetric responses to positive and negative information as seen in Section 5.3.

There is no significant effect on overall expenditures on education (either an effect on the gradient of investments or an average effect), attendance, or chores.

For transfers between schools, there is no significant change in the gradient of the investment function (cols. (1) - (4)), but, surprisingly, there is an ATE: treatment parents transferred their children to a different school 3 percentage points more than control parents (significant at the 1% level). Heterogeneity in the desired investment function by school type could explain these results: at schools with low average achievement, finding out a child is doing well might make it worth the transport or monetary costs of changing him to a better school, so transfers would be positively sloped with achievement. In contrast, at high-quality schools, finding out a child is doing poorly might be indicative of a poor match, and so the investment function would have the opposite slope. And, indeed, if we look at the results separately by school quality, proxied by school-average achievement, there are significant changes in the slopes of the investment functions, with the slope becoming more positive at low-quality schools and more negative at high-quality schools (Fig. 9a and Online App. Table A.5).

For parents' non-monetary investments, there is again no effect on the gradient of the investment function, but a positive ATE: a standardized index measure of non-monetary

investments increases by a statistically significant 0.65 standard deviations. Finding an ATE but no effect on the gradient could result from the investment index mixing complements with substitutes.[44] The result could also reflect heterogeneity in the investment function across the population, an uncertainty or salience effect, a Hawthorne effect (although both the treatment and control parents were aware that they were in an education study), or insufficient statistical power to detect the change in slope.

Uncertainty

As mentioned, one reason that we might not detect positive average treatment effects for dropouts even though dropouts increase with student achievement and parents were initially overconfident could be uncertainty. Column (4) of Appendix Table 3 tests for uncertainty effects by testing for treatment effects among parents whose baseline beliefs were close to the truth; unfortunately, power is too limited to be conclusive.

Detailed skills treatment results

Appendix Table 6 (cols. (2) and (3)) tests for whether the impacts of the information treatment are heterogeneous for parents who received the detailed skills report card relative to those who received the standard report card only, and finds that they are not.

7.3 Heterogeneity in the treatment effects by parent education

For the long-run investments, there is interesting heterogeneity by parent education that points to heterogeneity in how parents want to target their investments (See columns (1) - (4) of Table 7; Figure 9b shows the results graphically for expenditures and attendance). For expenditures, less-educated parents in the treatment group increase their spending on their lower-achieving children relative to their higher-achieving children. But, the more educated parents become, the more information causes them to spend relatively more on their higher-achieving children, until the gradient changes direction at roughly 5 years of education and treatment parents begin to spend more on their higher achievers.[45] A similar pattern holds for attendance: for the households of less-educated parents, information increases the attendance of low-achieving children, whereas it does the opposite for the households of more-educated parents. For both expenditures and attendance, the results are driven by investments on students who are in school, as the results are the same when one controls for

[44] The index represents the average standardized effect across all investments measured, where all are normalized so that positive means an increase in investment. Appendix Table 7 has the detailed treatment effects for each item in the index.

[45] Change in gradient figure based on linear extrapolation. The finding does not just result from the linear specification, however, as the conclusion is similar when one estimates the relationship separately within different education bins. These results are also robust to trimming the outliers, e.g., top-coding the data at the 99th percentile.

or conditions on enrollment. Figure 9b shows the attendance and expenditure heterogeneity results graphically.

In Section 6, the heterogeneity by parent education in the treatment effects was caused by heterogeneity in belief accuracy by parent education. Here, that factor may also play some role, but cannot explain the change in the sign of the effects. Rather, the results here suggest that there is heterogeneity in the production function and desired investment function. Less-educated parents might not expect that their children will be able to attend secondary school, and so could think there is a high return to attaining basic skills like reading. As a result, their desired investment functions would be negatively sloped. In contrast, more-educated parents might think their children have a chance of going to secondary school, and so could have high perceived returns from pushing high achievers over the hump into secondary school admissions, thereby causing a positively-sloped desired investment function. (Note that this story is more about wealth than education but education is the best-measured proxy for wealth in the data.) There are of course other possible explanations for the finding. The results do beg the question of why, even for the poor, information increased dropouts for the lower achieving students. It could be that some students are so far behind in achievement that they would need an infeasibly high level of expenditures to catch up, and so they drop out.

As outlined in Section 2, heterogeneity by parent education in the treatment effect on the slope of the investment function only shows us whether the investments of more- or less-educated parents are more distorted at baseline if there is no heterogeneity in the desired investment function. Thus, the results in Section 6.3, which used outcomes for which the desired investment functions are more homogeneous across SES levels, are better for understanding the relative level of distortions. The results presented in this section are more instructive about the differences in the desired investment functions.

Table 7 (cols. (5) and (6)) show that there is no statistically significant heterogeneity by parent education in average treatment effects for any investments examined; the predictions about heterogeneity were about reallocations, not levels.

8 Conclusion

This paper tests whether inaccuracies in parents' perceptions about their children's achievement impact investments in children's education. I find that there are large discrepancies between parents' beliefs about their children's achievement and their children's true achievement. At baseline, parents try to tailor their investments to their children's achievement levels, but their inaccurate beliefs prevent them from doing so. Providing parents with information significantly impacts their investments, allowing them to invest in the way that

they were trying to without information. Even within the fairly homogeneous, low-education context of Malawi, there is significant heterogeneity by parent education: Less-educated parents have less accurate beliefs, and update their beliefs and investments more in response to improved information. The heterogeneity in belief accuracy is also seen in other contexts, such as the U.S.. Taken together, the findings suggest that inaccurate beliefs may serve to perpetuate inequalities across generations within countries, as well as to provide one channel to explain why human capital levels across countries do not converge. The findings thus relate to a large literature on inter-generational mobility, both in developing countries (Hertz et al., 2007) and developed (Black and Devereux, 2011). They also advance our understanding of the role of misinformation in decision-making, relating to the literature not just in education but in other domains, such as health (Dupas, 2011; Madajewicz et al., 2007).

This paper also examines how parents' investments depend on their children's academic ability and endowments. Understanding this relationship is important for predicting policy spillovers: if parents spend more on their high-ability children, then policies that improve children's ability will crowd-in other household spending. The results here suggest that policies increasing ability may crowd-in spending for the rich but crowd-out spending for the poor, thereby increasing inequality.

It is perhaps surprising that baseline information is poor if the returns to knowledge are high and the information exists. However, parents may over-estimate their own knowledge, or (perceived) information acquisition costs may be high, both of which have been suggested in the U.S. in Bergman (2014). Qualitative interviews with parents also suggest that uneducated parents are too intimidated to talk to their children's teachers. This is consistent with the other findings in the literature that information constraints matter for education (Jensen, 2010; Dinkelman and Martínez A, 2014).

In general, an intervention that corrects one market imperfection can move us farther from the optimum if there are multiple market failures. One concern in this setting is that information may cause investments to decrease for some students. I do not find evidence of average decreases in investments, but there are decreases in certain investments for certain types of students (e.g., dropouts increase for lower-achieving students). If there are no other market failures, then this would improve social welfare, but there may be other reasons we think households underinvest in education for those students, such as agency issues within the household, externalities, or underestimation of returns. Policy makers may be concerned about these issues for a scale-up, although the first-best solution would be to deliver the information and use other policies to directly target the other market failures, since withholding information has obvious costs (e.g., efficiency).

Another potential concern for a scale-up would be if beliefs enter directly into the utility

function, an assumption which is sometimes made in the behavioral economics literature (e.g., Köszegi (2006), Bénabou and Tirole (2002), and Weinberg (2009)) and can imply that some overconfidence is optimal. If beliefs have motivational value, as suggested in Bénabou and Tirole (2002), then increasing information could have earnings or wealth costs, although we might expect this channel to be less relevant when talking about parents' confidence, not own confidence.

This paper is focused on identifying the causal chain between perceptions and investments, not on designing a cost-effective information policy. The policy design issues (such as whether information delivery through schools can be improved, how frequently information should be delivered) are left open for future research.

References

Aizer, A. and F. Cunha (2012). The production of child human capital: Endowments, investments and fertility. *Unpublished paper, Brown University*.

Alexander, K. and D. Entwisle (2006). Baltimore beginning school study, 1982-2002. The Harvard-MIT Data Centers. Henry A. Murray Research Archive. Log# 01293.

Andrabi, T., J. Das, and A. I. Khwaja (2009). Report cards: The impact of providing school and child test scores on educational markets. *Unpublished Working Paper*.

Andrabi, T., J. Das, and A. I. Khwaja (2012). What did you do all day? maternal education and child outcomes. *Journal of Human Resources 47*(4), 873–912.

Banerjee, A. V., R. Banerji, E. Duflo, R. Glennerster, and S. Khemani (2010). Pitfalls of participatory programs: Evidence from a randomized evaluation in education in india. *American Economic Journal: Economic Policy*, 1–30.

Banerji, R., J. Berry, and M. Shotland (2013). The impact of mother literacy and participation programs on child learning: evidence from a randomized evaluation in india. *Cambridge, MA: Abdul Latif Jameel Poverty Action Lab (J-PAL)*.

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy 70*(5), 9–49.

Behrman, J. R., M. R. Rosenzweig, and P. Taubman (1994). Endowments and the allocation of schooling in the family and in the marriage market: The twins experiment. *Journal of Political Economy*, 1131–1174.

Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 871–915.

Bergman, P. (2014). Parent-child information frictions and human capital investment: Evidence from a field experiment. *Unpublished Working paper*.

Black, S. E. and P. J. Devereux (2011). Recent developments in intergenerational mobility. *Handbook of labor economics 4*, 1487–1541.

Chevalier, A., S. Gibbons, A. Thorpe, M. Snell, and S. Hoskins (2009). Students' academic self-perception. *Economics of Education Review 28*(6), 716–727.

Datar, A., M. R. Kilburn, and D. S. Loughran (2010). Endowments and parental investments in infancy and early childhood. *Demography 47*(1), 145–162.

Dinkelman, T. and C. Martínez A (2014). Investing in schooling in Chile: The role of information about financial aid for higher education. *Review of Economics and Statistics 96*(2), 244–257.

Dupas, P. (2011). Do teenagers respond to hiv risk information? evidence from a field experiment in kenya. *American Economic Journal: Applied Economics*, 1–34.

Griliches, Z. (1979). Sibling model and data in econometrics: Beginnings of a survey. *Journal of Political Economy 87*, S37–S64.

Hertz, T., T. Jayasundera, P. Piraino, S. Selcuk, N. Smith, and A. Verashchagina (2007). The inheritance of educational inequality: International comparisons and fifty-year trends. *The BE Journal of Economic Analysis & Policy 7*(2).

Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics 125*(2), 515–548.

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association 4*(4), 673–707.

Leight, J. (2014). Sibling rivalry: Endowment and intrahousehold allocation in gansu province, china.

Madajewicz, M., A. Pfaff, A. Van Geen, J. Graziano, I. Hussein, H. Momotaj, R. Sylvi, and H. Ahsan (2007). Can information alone change behavior? response to arsenic contamination of groundwater in bangladesh. *Journal of Development Economics 84*(2), 731–754.

Nguyen, T. (2008). Information, role models and perceived returns to education: Experimental evidence from madagascar. *Unpublished manuscript*.

Oreopoulos, P., M. E. Page, and A. H. Stevens (2006). The intergenerational effects of compulsory schooling. *Journal of Labor Economics 24*(4), 729–760.

Paxson, C. H. and N. R. Schady (2005). *Cognitive development among young children in Ecuador: The roles of wealth, health and parenting*, Volume 3605. World Bank Publications.

Pitt, M. M., M. R. Rosenzweig, and M. N. Hassan (1990). Productivity, health, and inequality in the intrahousehold distribution of food in low-income countries. *The American Economic Review*, 1139–1156.

Rosenzweig, M. R. and K. I. Wolpin (1994). Are there increasing returns to the intergenerational production of human capital? maternal schooling and child intellectual achievement. *Journal of Human Resources*, 670–693.

Rosenzweig, M. R. and J. Zhang (2009). Do population control policies induce more human capital investment? twins, birth weight and china's one-child policy. *The Review of Economic Studies 76*(3), 1149–1174.

Stinebrickner, T. R. and R. Stinebrickner (2009). Learning about academic ability and the college drop-out decision. Technical report, National Bureau of Economic Research.

Weinberg, B. A. (2009). A model of overconfidence. *Pacific Economic Review 14*(4), 502–515.

Figure 1: Motivating evidence

Less-educated parents have less accurate perceptions about their children's academic performance than more-educated parents

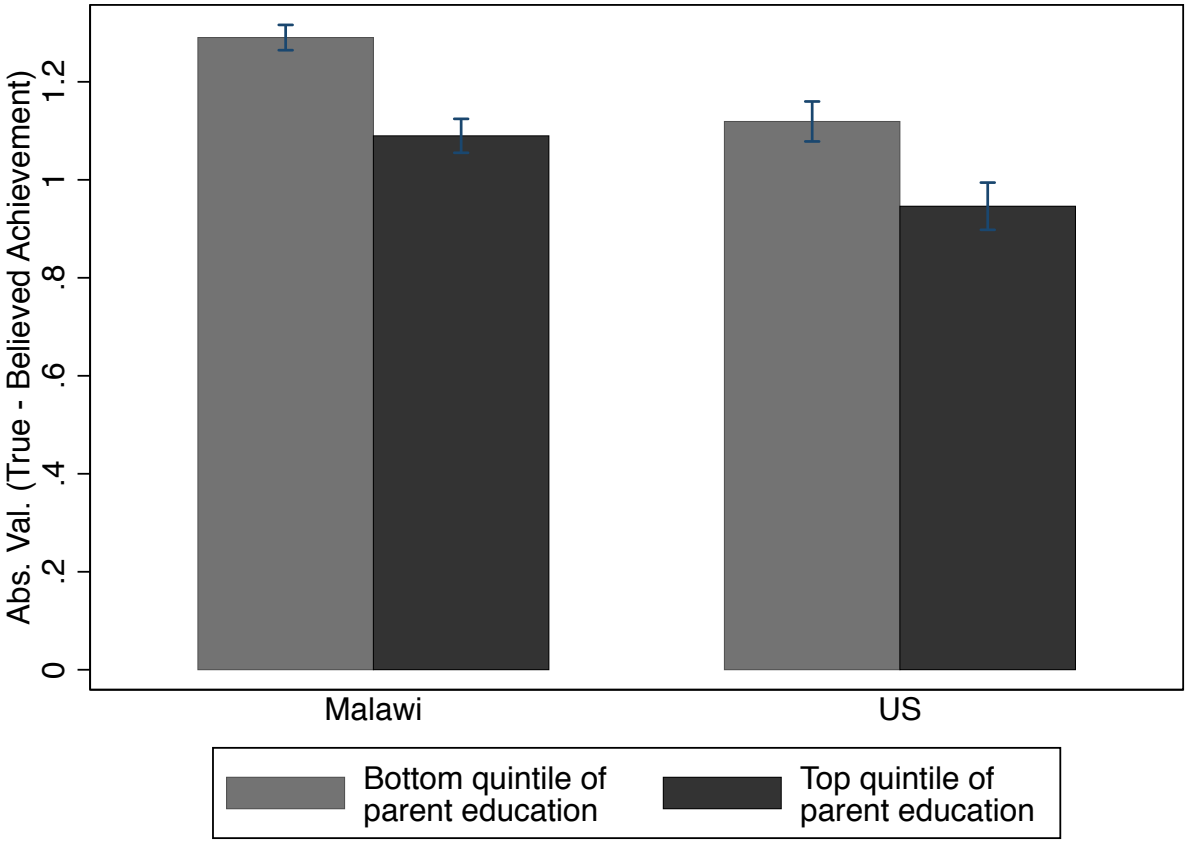

Notes: Bar height represents the absolute value of the gap between parents' perceptions about their child's academic performance and their child's true academic performance, expressed in standard deviations of the academic performance distribution. The U.S. data source is the Beginning School Study (Alexander and Entwisle, 2006), a longitudinal study of children's academic and social development conducted with families of children enrolled in Baltimore City Public Schools. The academic performance measure is student performance (grades) on school report cards. The Malawi data are from data collected for this study, and the academic performance measure is student achievement from tests conducted in schools. Results robust to using other measures of education (e.g., above-median). In Malawi, the bottom quintile is 1 year of education or fewer, the top quintile is 8+, and the median is 4. In the U.S., the bottom quintile is 10 years or fewer, the top quintile is 13.5+, and the median is 12.

Figure 2: Theoretical predictions

Inaccurate perceptions could cause investment gradient on truth to be flatter than on beliefs

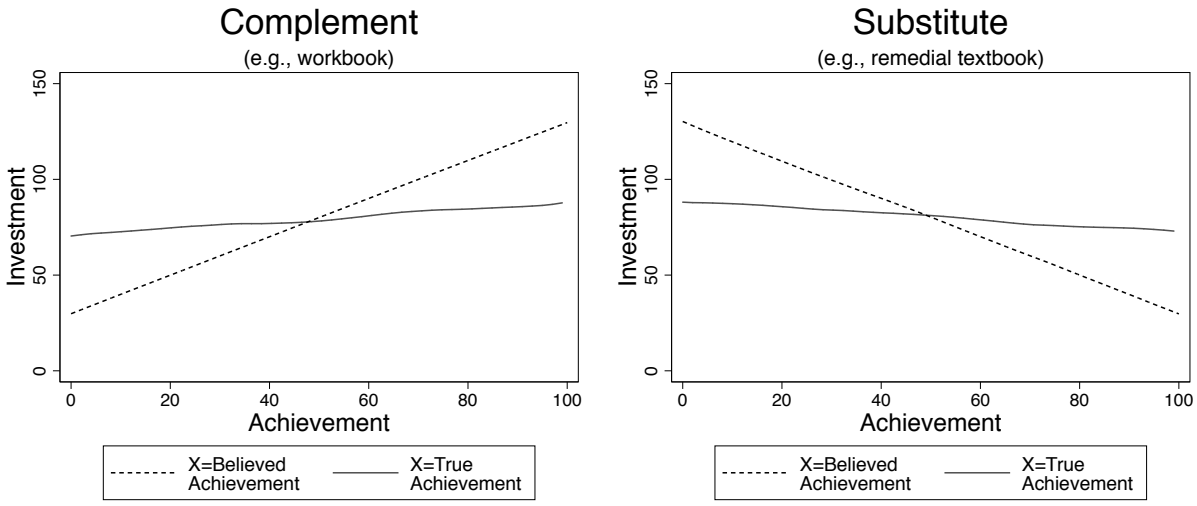

Notes: Illustrative graph, not based on real data. The two lines have the same y-axis but different x-axes: the dashed line has believed achievement as the x-axis, whereas the solid line has true achievement as the x-axis.

Figure 3: Parents have inaccurate perceptions about their children's achievement

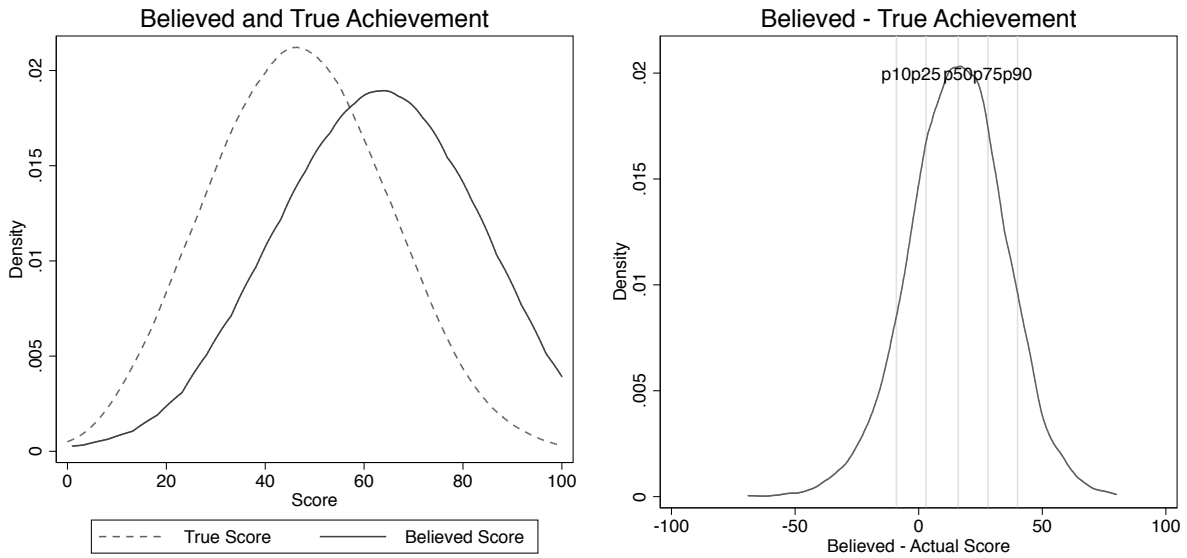

Notes: Data source is baseline data (full sample). The left graph shows kernel density plots comparing the distribution of parents' beliefs about their children's Term 2 2011-2012 achievement test performance, elicited at the beginning of the baseline survey, with the distribution of their children's true Term 2 achievement test performance. The right graph shows a kernel density plot of the distribution, across parents, of each parent's beliefs about her child's achievement relative to her child's true achievement. The lines represent the percentiles of the distribution.

Figure 4: Consistent with a distortion, in the control group, the investment gradient on true achievement is flatter than on believed achievement

(a) **Workbooks (Complements)**

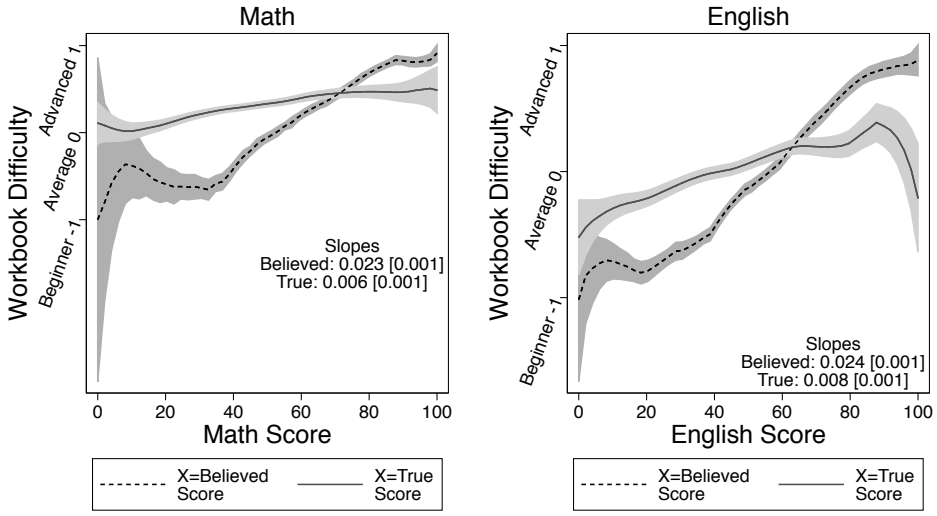

(b) **Textbook WTP (Substitute)**

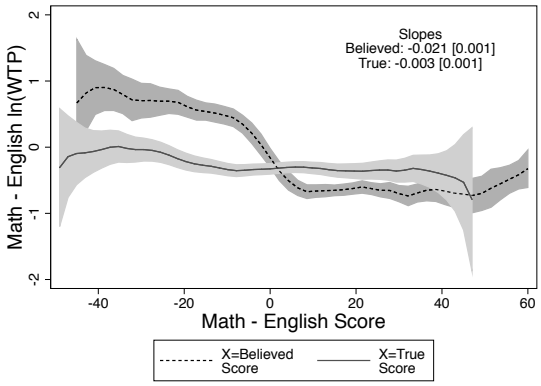

(c) **Secondary School Lottery**

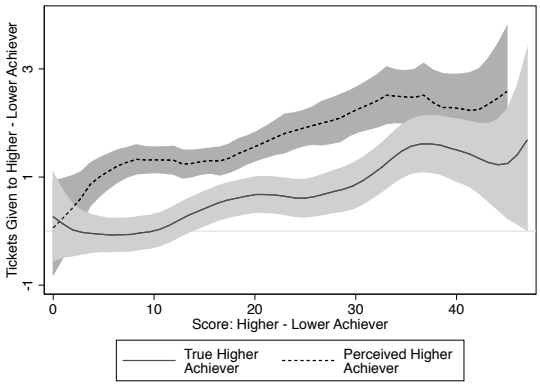

Notes: Data source is baseline data from the control group only. All lines are locally linear regression lines with investments as the dependent variable and either true (solid line) or perceived (dashed line) achievement as the x-axis. Beliefs were elicited from parents at the beginning of the baseline survey. For the workbook graphs (panel (a)), the dependent variable is the parent's choice among 3 level-specific workbooks which are parametrized as -1 (beginner), 0 (average) and 1 (advanced). For textbook WTP (panel (b)), the dependent variable is the difference in the parent's log WTP for a remedial math textbook relative to a remedial English textbook. For the secondary school lottery, the dependent variable is the number of secondary school lottery tickets given to the higher relative to the lower achiever. So, for the solid line, the dependent variable is the number of tickets given to the higher achiever relative to the lower achiever and the x-axis is the true achievement gap (higher - lower achiever). For the dashed line, the dependent variable is the number of tickets given to the perceived higher achiever relative to the perceived lower achiever and the x-axis is the perceived achievement gap. The grey areas represent 95% confidence intervals.

Figure 5: Information shifts parents' beliefs towards their children's true achievement

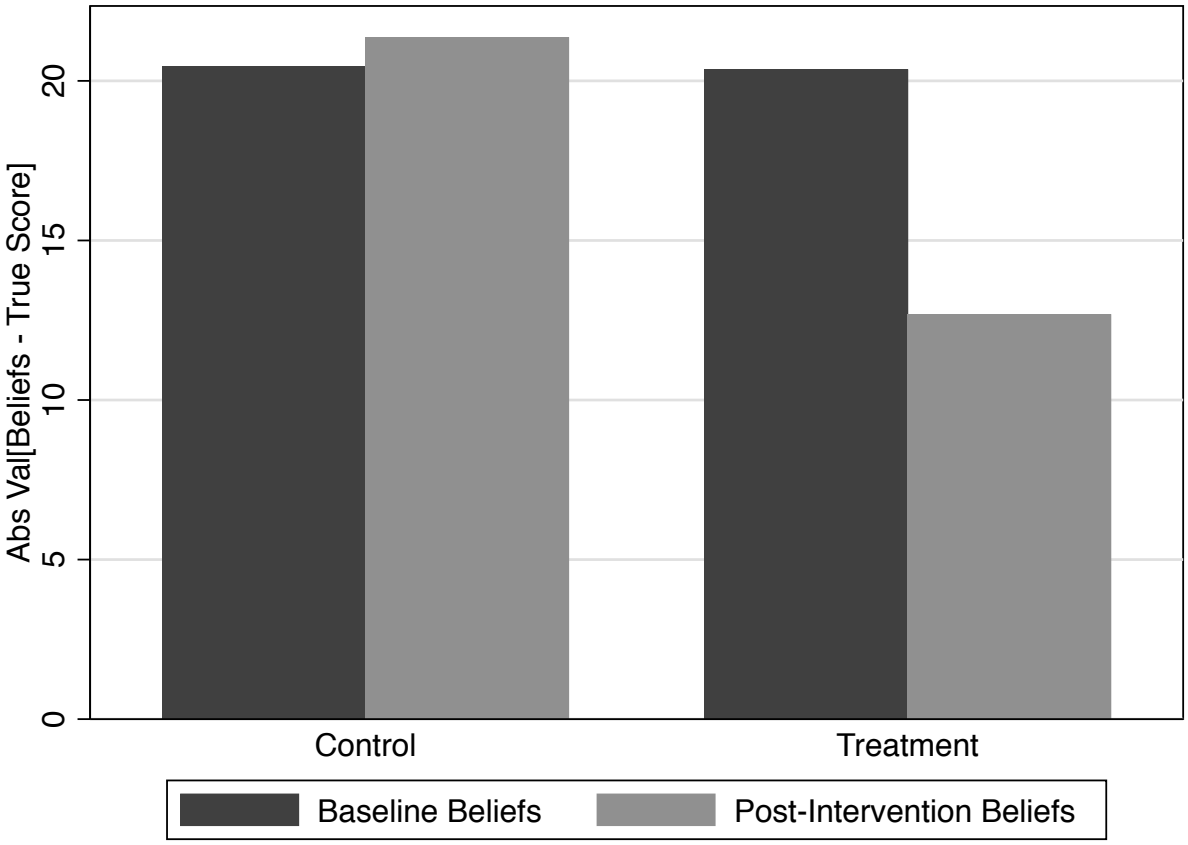

Notes: The dark gray bars show the absolute value of the difference between children's true term 2 2011-2012 achievement test scores and their parents' baseline beliefs about those scores, which were elicited at the beginning of the baseline survey (before the information treatment). The light gray bars show the absolute value of the difference between children's true term 2 achievement test scores and their parents' beliefs, elicited after the intervention, about their children's hypothetical scores if they took an achievement test on that same day. The p-value for equality between the treatment and control groups for the height of the dark grey bars is .825 (i.e., there is balance) while the p-value for equality between the treatment and control groups for the height of the light grey bars is $< .01$, as is the p-value for the difference between the heights of the dark and light gray bars for the treatment group.

Figure 6: Treatment effects: Information increases the gradient of investments on true achievement

(a) **Workbooks (Complements)**

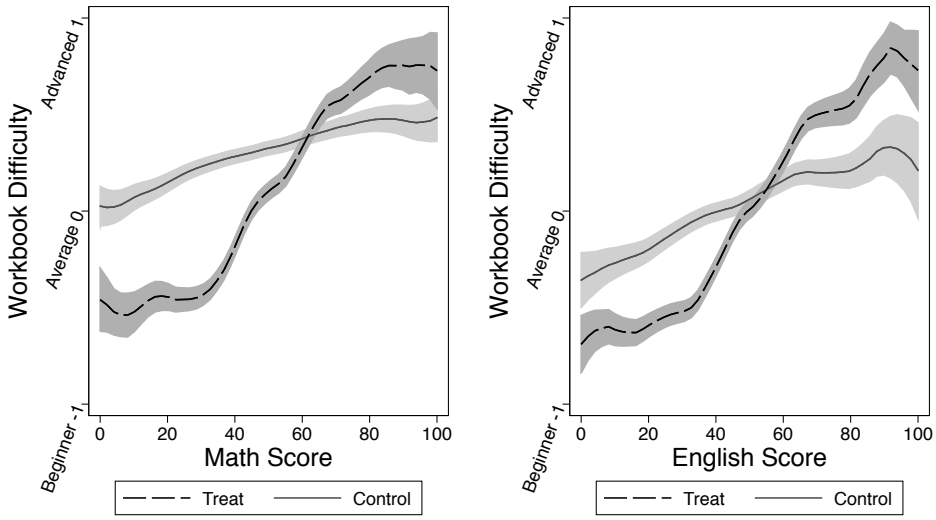

(b) **Textbook WTP (Substitute)**

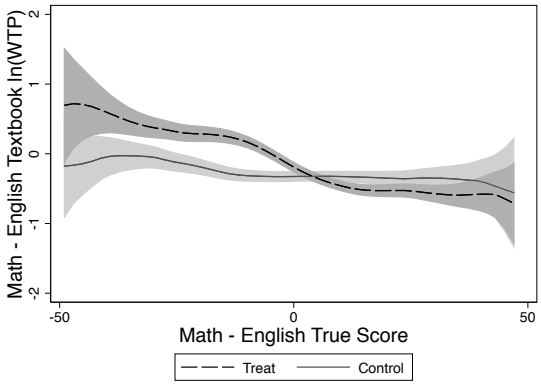

(c) **Secondary School Lottery**

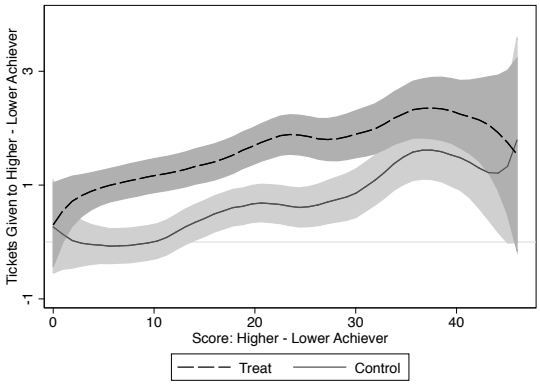

Notes: All lines are locally linear regression lines with investments as the dependent variable and true achievement as the x-axis. For the workbook graphs, the dependent variable is the parent's choice among 3 level-specific workbooks which are parametrized as -1 (beginner), 0 (average) and 1 (advanced). For textbooks, the dependent variable is the difference in the parent's log WTP for a remedial math textbook relative to a remedial English textbook. For the secondary school lottery, the dependent variable is tickets given to the higher relative to the lower achiever. The grey areas represent 95% confidence intervals.

Figure 7: Heterogeneity in treatment effects by parent education (Textbooks, workbooks, secondary school lottery)

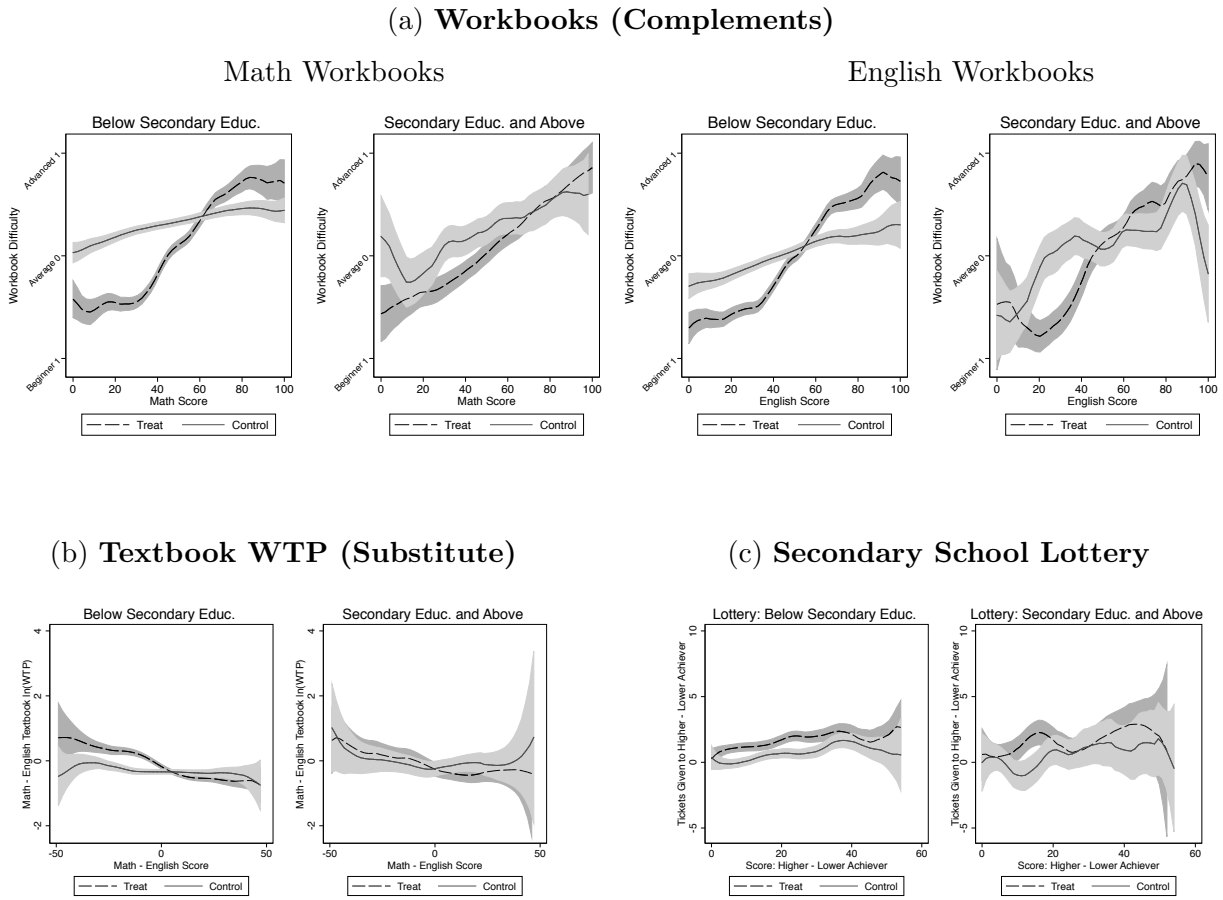

Notes: Left graphs for each outcome are estimated for respondents without secondary education, right graphs are estimated for respondents with any secondary education and above. All lines are locally linear regression lines with investments as the dependent variable and true achievement as the x-axis. For the workbook graphs, the dependent variable is the parent's choice among 3 level-specific workbooks which are parametrized as -1 (beginner), 0 (average) and 1 (advanced). For textbooks, the dependent variable is the difference in the parent's log WTP for a remedial math textbook relative to a remedial English textbook. For the secondary school lottery, the dependent variable is tickets given to the higher relative to the lower achiever. The grey areas represent 95% confidence intervals.

Figure 8: Treatment effects: Medium-run and longer-run outcomes

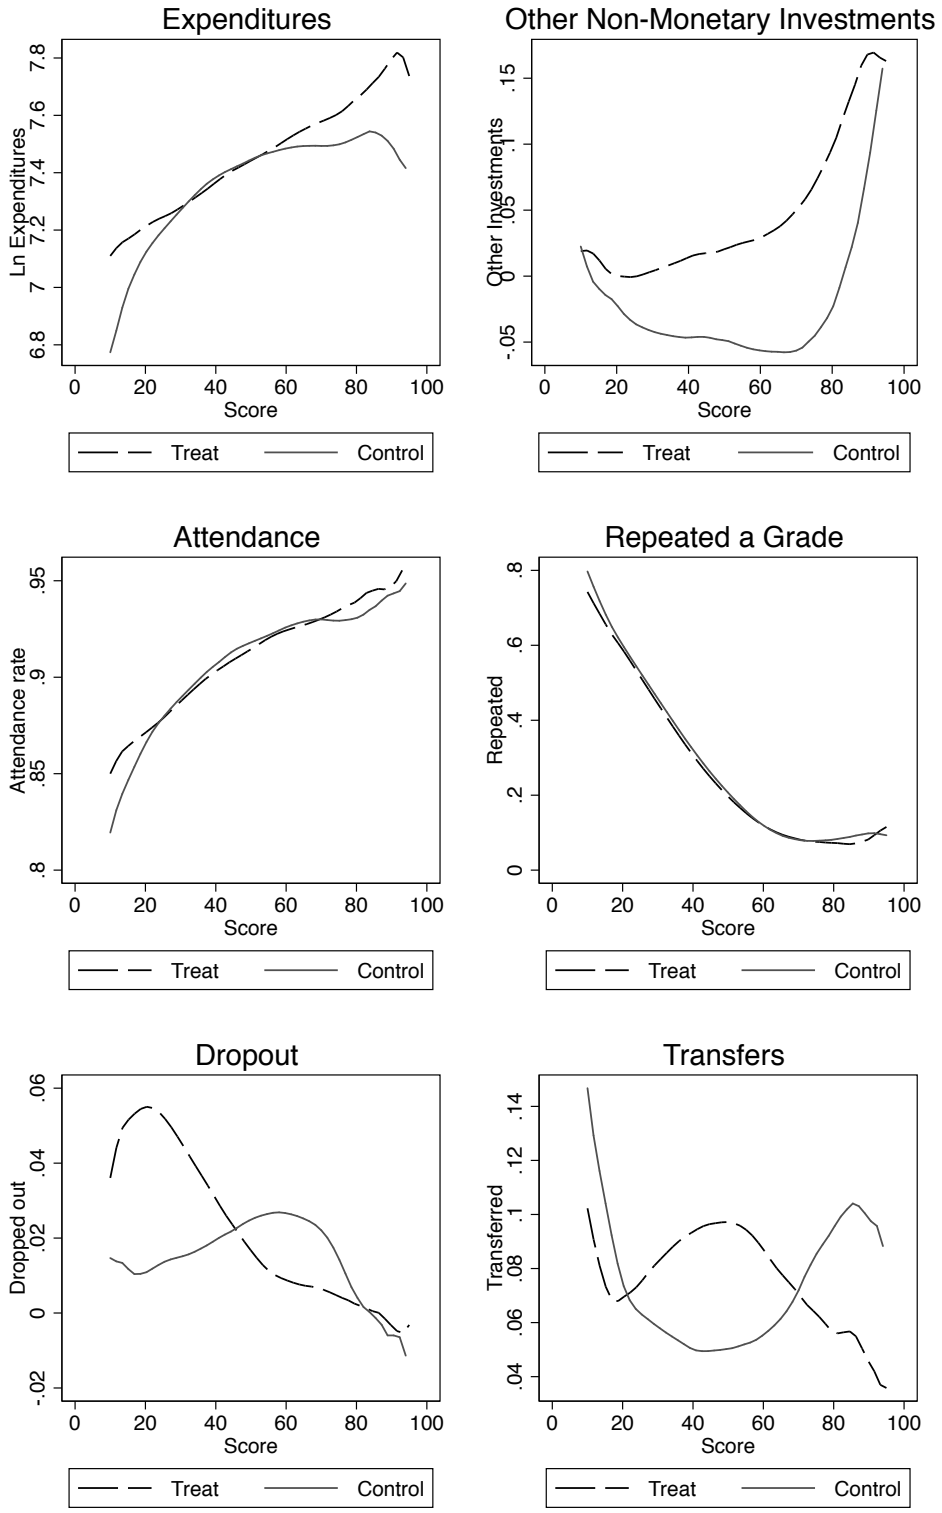

Notes: Data sources are endline survey data (expenditures, other non-monetary investments, homework help, asked for homework help, dropout, and transfers) and data collected from schools (attendance rate, grade repetition). All lines are locally linear regression lines with investments as the dependent variable and true baseline achievement as the x-axis.

Figure 9: Heterogeneity in longer-run treatment effects (Selected outcomes)

(a) **Heterogeneity by school average achievement**

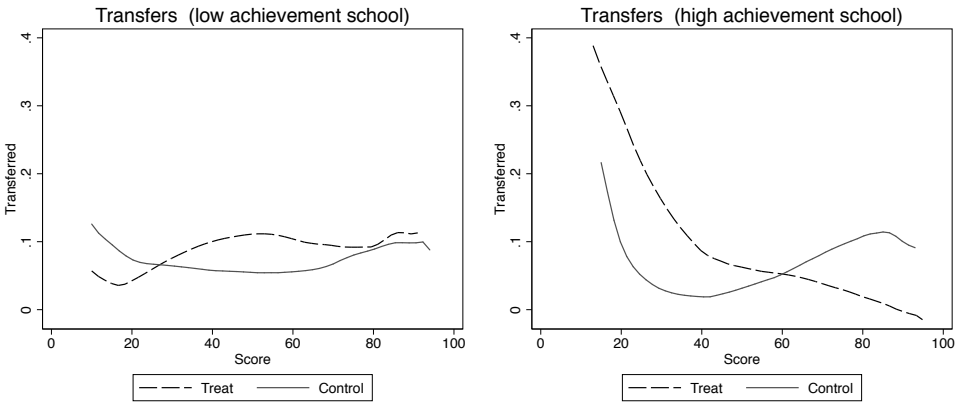

(b) **Heterogeneity by parent education**

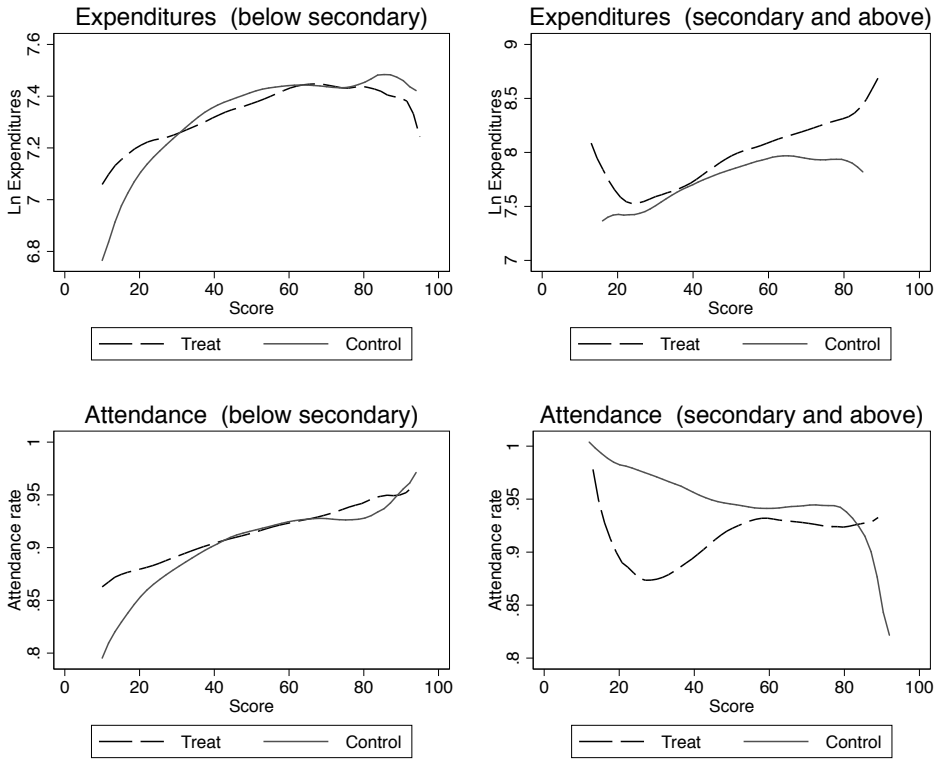

Notes: The figure contains the same figures displayed in Fig. 8 but estimated separately for different subsamples (see notes for Fig. 8 for more detailed description). In panel (a), the results are estimated separately for schools in the top quartile of overall student achievement (right graph) and schools not in the top quartile (left graph). In panel (b), the results are estimated separately for respondents without secondary education (left column) and with any secondary education and above (right column).

Table 1: Baseline Summary Statistics

	Full Sample		Control	Treat	Treat − Control		
	Mean	SD	Mean	Mean	Mean	Std. Error	p-val T=C
Respondent Background							
Female	0.77	0.42	0.77	0.76	-0.01	0.02	0.37
Primary education decision maker	0.92	0.27	0.91	0.92	0.01	0.01	0.31
Age	40.8	11.0	40.6	41.0	0.32	0.44	0.47
Education (years)	4.44	3.57	4.42	4.45	0.04	0.13	0.78
Respondent has secondary education +	0.11	0.31	0.11	0.11	0.01	0.01	0.62
Parent can read or write Chichewa	0.67	0.47	0.67	0.68	0.01	0.02	0.67
Respondent is farmer	0.46	0.5	0.47	0.46	-0.01	0.02	0.7
Respondent's weekly income	2,126	4,744	2,051	2,203	197	194	0.31
Household Background							
Number of children[a]	5.13	1.74	5.16	5.1	-0.05	0.07	0.47
One-parent household	0.19	0.39	0.19	0.2	0.01	0.02	0.47
Parents' average education (years)	4.66	3.25	4.68	4.64	-0.04	0.12	0.74
Any parent has secondary education +	0.18	0.38	0.17	0.19	0.02	0.01	0.24
Student Information							
Child's grade level	3.72	1.37	3.72	3.72	0	0.04	0.94
Child's age	11.6	2.68	11.7	11.6	-0.1	0.08	0.21
Child is female	0.51	0.5	0.52	0.5	-0.02	0.01	0.25
Baseline attendance	0.91	0.13	0.92	0.91	0	0	0.72
Annual per-child education expenditures	1,742	2,791	1,712	1,772	58.0	83.0	0.48
Fees paid to schools	381	1,128	384	378	-6.84	23.9	0.78
Uniform expense	576	1,019	548	603	49.9	36.1	0.17
School supplies, books, tutoring, etc.	785	1,819	780	790	14.3	62.3	0.82
Any supplementary expenditures on child	0.9	0.3	0.9	0.89	-0.01	0.01	0.49
Achievement Scores							
Overall score	46.8	17.5	47.1	46.4	-0.74	0.46	0.11
Math score	44.9	20.2	45.4	44.4	-1.08	0.54	0.04
English score	44.2	20.1	44.5	43.9	-0.56	0.53	0.29
Chichewa score	51.3	22.6	51.5	51.0	-0.57	0.59	0.34
(Math − English) Score	0.71	19.5	0.93	0.5	-0.53	0.51	0.3
Respondent's Beliefs about Child's Achievement Scores							
Believed Overall Score	62.4	16.5	62.7	62.0	-0.78	0.48	0.11
Believed Math Score	64.7	19.0	65.2	64.3	-0.94	0.55	0.09
Believed English Score	55.3	20.9	55.6	54.9	-0.71	0.62	0.25
Believed Chichewa Score	66.8	19.4	66.8	66.7	-0.1	0.6	0.87
Beliefs about (Math − English) Score	9.48	21.5	9.59	9.37	-0.23	0.63	0.71
Respondent's Misperception about Child's Achievement							
Abs Val [Believed − True Overall Score]	20.4	14.5	20.4	20.3	-0.12	0.43	0.77
Abs Val [Believed − True Math Score]	25.8	18.0	25.8	25.7	-0.1	0.52	0.85
Abs Val [Believed − True English Score]	21.4	16.4	21.6	21.1	-0.57	0.48	0.23
Abs Val [Believed − True Chichewa Score]	23.8	17.5	23.7	23.9	0.18	0.51	0.73
Abs Val [Believed − True (Math-English) Score]	22.1	17.4	22.3	21.9	-0.44	0.51	0.39
Abs Val [Believed − True Overall Score (Child1-2)]	18.7	15.1	18.9	18.5	-0.35	0.59	0.55
Beliefs about Returns to Education							
Returns to educ. (sec. school/prim. earnings)	3.22	3.79	3.28	3.16	-0.11	0.15	0.47
Believes educ. and achievement complementary	0.91	0.29	0.9	0.91	0	0.01	0.68
Sample Sizes							
Sample Size–HHs	2,634		1,327	1,307			
Sample Size–Kids	5,268		2,654	2,614			

Notes. Data Source is baseline survey. Standard errors for the t-test of equality clustered at the household level.
a. Counted as a child if either of the primary caregivers for the sampled children is a parent of the child.
b. Includes exercise books and pencils, textbooks and supplementary reading books, backpacks, and tutoring expenses.
c. Respondent said that they thought the earnings of a more able child would increase "more" or "much more" than the earnings of a less able child from getting a secondary education.

Table 2: Information treatment effects (Textbooks, workbooks, and sec. school fee lottery)

	Treatment effect on slope					A.T.E.
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Dependent Var: Math Workbook Choice						
Treat × Math Score	1.34***	1.20***	1.20***	1.13***	1.13***	
	[0.093]	[0.17]	[0.17]	[0.17]	[0.17]	
Treat	-91.1***					-31.0***
	[4.91]					[2.05]
Math Score	0.65***	0.81***	0.81***	0.84***	0.84***	1.32***
	[0.065]	[0.12]	[0.12]	[0.12]	[0.12]	[0.051]
Household FE		✓	✓	✓	✓	
Observations	5,239	5,239	5,239	5,239	5,239	5,239
R-squared	0.217	0.695	0.695	0.696	0.696	0.184
Panel B. Dependent Var: English Workbook Choice						
Treat × English Score	1.26***	1.27***	1.26***	1.33***	1.33***	
	[0.096]	[0.17]	[0.17]	[0.17]	[0.17]	
Treat	-68.5***					-13.0***
	[4.83]					[2.15]
English Score	0.76***	0.89***	0.89***	0.86***	0.85***	1.39***
	[0.073]	[0.12]	[0.12]	[0.12]	[0.12]	[0.052]
Household FE		✓	✓	✓	✓	
Observations	5,239	5,239	5,239	5,239	5,239	5,239
R-squared	0.204	0.710	0.710	0.714	0.715	0.177
Panel C. Dependent Var: ln(Math Textbook WTP) - ln(English Textbook WTP)						
Treat × (Math − English Score)	-0.013***	-0.013***	-0.013***	-0.014***	-0.014***	
	[0.0022]	[0.0038]	[0.0038]	[0.0039]	[0.0039]	
Treat	0.15***					0.14***
	[0.041]					[0.041]
Math − English Score	-0.0030*	-0.0016	-0.0015	-0.00048	-0.00041	-0.0099***
	[0.0016]	[0.0025]	[0.0025]	[0.0028]	[0.0028]	[0.0011]
Household FE		✓	✓	✓	✓	
Observations	5,183	5,183	5,183	5,183	5,183	5,183
R-squared	0.033	0.601	0.601	0.602	0.602	0.024
Panel D. Dependent Var: Lottery tickets received						
Treat × (Higher-scoring Sibling)	0.98***	0.98***	0.98***	0.94***	0.95***	
	[0.13]	[0.22]	[0.22]	[0.21]	[0.22]	
Treat × (Overall Score)		0.0017	0.0017	0.0052	0.0036	
		[0.0090]	[0.0090]	[0.0088]	[0.0091]	
Higher-scoring Sibling	0.53***	-0.16	-0.17	-0.16	-0.19	
	[0.091]	[0.16]	[0.15]	[0.15]	[0.16]	
Overall score		0.034***	0.034***	0.031***	0.033***	
		[0.0064]	[0.0064]	[0.0063]	[0.0064]	
Household FE	✓	✓	✓	✓	✓	
Observations	5,258	5,258	5,258	5,258	5,080	
R-squared	0.105	0.125	0.129	0.161	0.175	
Column Includes Controls for:						
Treat × Female			✓	✓	✓	
Treat × Grade Level				✓	✓	
Treat × Educ. Expenditures					✓	

Notes: Each observation is a child. Standard errors clustered at the household level. Regressions control for school FE, parents' education level, the between-child score gap, child achievement, and the main effect of any variables interacted with Treat. Workbook choices are -100 for beginner, 0 for average, 100 for advanced. The treatment effect on slope results can be interpreted as follows: Take for example Panel A., column (1). The coefficient on Math Score is the slope of the line in the control group: if a child's score increases by one point, the chance that her parent chooses the next level of workbook increases by .65%. The coefficient on Treat x Score represents the treatment effect on the slope; the coefficient of .013 means the treatment increased the slope by 200% (.013/.0065). A.T.E. stands for Avg. Treatment Effect. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Less educated parents have less accurate beliefs

	Dependent Variable= Abs Val [True − Believed Achievement]					
	(1)	(2)	(3)	(4)	(5)	(6)
Performance Measure:	Overall	Overall	Math	Math	English	English
Parents' years of education	-0.201***	-0.197***	-0.258***	-0.274***	-0.106	-0.105
	[0.066]	[0.066]	[0.080]	[0.079]	[0.074]	[0.074]
Observations	5,019	5,019	5,021	5,021	5,021	5,021
Dep var mean	20.41		25.82		21.42	
	(7)	(8)	(9)	(10)	(11)	(12)
Performance Measure:	Chichewa	Chichewa	Math−Eng	Math−Eng	Child 2−1	Child 2−1
Parents' years of education	-0.327***	-0.299***	-0.025	-0.035	-0.257***	-0.240***
	[0.078]	[0.077]	[0.078]	[0.078]	[0.091]	[0.091]
Observations	5,021	5,021	5,021	5,021	2,514	2,514
Dep var mean	23.84		22.12		18.73	
Col. Specification Details						
Child and Parent Controls		✓		✓		✓

Notes. Robust standard errors in brackets. Standard errors clustered at the household level. Child and parent controls include a control for child gender, grade FE, parent gender, and whether the parent is the primary education decisionmaker. Parents' years of education is the average across parents in the household.
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Less-educated parents update their individual beliefs more than more-educated parents do, but the mean of the beliefs distribution does not shift significantly more for less-educated parents

Dependent Variable:	Abs Val [Post Intervention Beliefs − Baseline]		Post Intervention Beliefs − Baseline	
	(1)	(2)	(3)	(4)
Treat × (Parent yrs. education)	-0.37***	-0.29***	0.19	-0.0033
	[0.11]	[0.098]	[0.14]	[0.13]
Treat	10.0***	-3.30***	-7.53***	-12.8***
	[0.64]	[1.26]	[0.84]	[2.05]
Parent yrs. education	-0.021	-0.019	0.015	0.0085
	[0.061]	[0.061]	[0.072]	[0.072]
Treat × Score Control		✓		✓
Treat × | Beliefs − Truth | Control		✓		✓
Observations	4,951	4,951	4,951	4,951
R-squared	0.133	0.305	0.053	0.233
Dep Var Mean in Treat	13.72		-5.91	
Dep Var Mean in Control	5.456		0.722	

Notes: Standard errors clustered at the household level. Baseline beliefs were elicited before the information intervention about Term 2 2011-2012 achievement (the same metric delivered to parents. Post-intervention beliefs were elicited after the information intervention about the child's achievement if they were to take an achivement test that day. Parent years of education is the average across parents in the household. Regressions control for school FE, parents' education level, the between-child score gap, and child achievement.
*** p<0.01, ** p<0.05, * p<0.1

Table 5: Treatment effect heterogeneity by parent education

	Math Workbook			**English Workbook**			**Textbook WTP**			**Lottery**		
Dependent Variable:	100 for Advanced, 0 for Average, -100 for Beginner			100 for Advanced, 0 for Average, -100 for Beginner			ln(Textbook WTP) for Math−English			Tickets given to higher−lower achiever		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Sample:	All	Treat Only	All	All	Treat Only	All	All	Treat Only	All	All	Treat Only	All
Treat × Score	1.92***			1.57***			-0.017***			0.0033		
	[0.16]			[0.17]			[0.0037]			[0.017]		
Treat × Score × (Parent yrs of educ.)	-0.12***			-0.066**			0.00073			-0.00041		
	[0.027]			[0.029]			[0.00059]			[0.0029]		
Score	0.29**	2.14***		0.61***	2.17***		0.000055	-0.017***		0.033***	0.038***	
	[0.11]	[0.12]		[0.13]	[0.11]		[0.0026]	[0.0027]		[0.012]	[0.013]	
Score × (Parent yrs of educ.)	0.078***	-0.041**		0.032	-0.032*		-0.00058	0.00014		0.00047	0.000064	
	[0.020]	[0.019]		[0.022]	[0.019]		[0.00038]	[0.00045]		[0.0020]	[0.0020]	
Treat	-121.5***		-35.3***	-79.1***		-12.3***	0.30***		0.27***	1.05***		1.12***
	[8.58]		[3.64]	[8.59]		[3.82]	[0.071]		[0.071]	[0.40]		[0.23]
Treat × (Parent yrs of educ.)	6.48***		0.93	2.29		-0.13	-0.032***		-0.027**	-0.015		-0.023
	[1.46]		[0.61]	[1.53]		[0.65]	[0.012]		[0.012]	[0.069]		[0.040]
Parent yrs of educ.	-3.86***	2.34**	-0.22	-0.29	1.88*	0.97*	0.024***	-0.0083	0.023***	-0.00044	-0.016	0.0090
	[1.08]	[1.04]	[0.47]	[1.18]	[0.97]	[0.53]	[0.0084]	[0.0089]	[0.0084]	[0.050]	[0.047]	[0.029]
Score Used	Math	Math	Math	English	English	English	Math − English	Math − English	Math − English	Higher − Lower Child	Higher − Lower Child	Higher − Lower Child
Observations	5,203	2,588	5,203	5,203	2,588	5,203	5,183	2,575	5,183	2,611	1,299	2,611
R-squared	0.220	0.292	0.184	0.207	0.325	0.179	0.035	0.059	0.014	0.047	0.028	0.047
p-val: Treat × Perf × Educ=0	5.0e-06			0.022			0.222			0.888		

Notes. Robust standard errors in brackets. Standard errors clustered at household level. The first column for each outcome variable shows the heterogeneity by parent education in the information treatment effect on the gradient of the investment function. The second column shows the heterogeneity by parent education in the gradient of the investment function in the treatment group. The third column shows the heterogeneity in the average effect of the treatment. Each observation is a child (cols. (1)-(9)) or a household (cols. (10)-(12)). Parent years of education is the average across parents in the household. Regressions control for school FE, parents' education level, the between-child score gap, and child achievement.
*** p<0.01, ** p<0.05, * p<0.1

Table 6: Information treatment effects: Longer-run outcomes

	Coefficient estimate (standard error) for:						
	A. Heterogeneity in treatment effects by performance – Linear spec		B. Heterogeneity in treatment effects by performance – Above-median spec		C. Ave. treat-ment effect		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Independent Variable:	Treat	Treat× Score	Treat	Treat× Above-Median Score	Treat	Control group mean	N
Dependent Variables							
Panel A. Dropout and transfer (from endline survey data)							
Dropout	0.055	-0.0011	0.022	-0.037	0.004	0.021	1,786
	[0.021]***	[0.0004]***	[0.012]*	[0.015]***	[0.007]		
Transfer	0.023	0.0002	0.022	0.017	0.03	0.057	1,781
	[0.036]	[0.0007]	[0.019]	[0.025]	[0.014]**		
Panel B. Investments (from endline survey data)							
Total educ. expenditures	119.70	-0.325	100.54	4.181	104.45	2,362.06	1,729
	[291.50]	[6.841]	[177.56]	[230.00]	[164.32]		
ln(Total educ. expenditures)	0.093	-0.0019	0.014	-0.03	0.0013	7.389	1,709
	[0.114]	[0.002]	[0.061]	[0.074]	[0.049]		
Avg. std. effect across other non-monetary investments[a,b]	0.07	-0.0001	0.057	0.015	0.065	-0.012	1,720
	[0.057]	[0.0011]	[0.032]*	[0.039]	[0.026]***		
Avg. std. effect across other chores[c]	0.01	0.001	0.034	0.049	0.058	-0.0009	1,681
	[0.104]	[0.002]	[0.05]	[0.069]	[0.041]		
Panel C. Attendance and grades (from data collected from schools)							
Attendance rate following baseline survey	-0.008	0.0001	-0.0015	-0.0017	-0.002	0.911	1,827
	[0.026]	[0.0005]	[0.012]	[0.015]	[0.008]		
End-of-year grade	0.122	-0.003	0.03	-0.095	-0.016	1.97	1,241
	[0.091]	[0.0019]	[0.047]	[0.07]	[0.036]		

Notes. Data sources are endline survey and endline data collected from schools. Each observation is a child. Standard errors clustered at the household level. All regressions control for child baseline achievement, school fixed effects, parents' education, the between-sibling achievement gap, grade fixed effects, and the baseline value of the dependent variable, if available (not available for dropouts, transfers, pushing children to attend school). *** p<0.01, ** p<0.05, * p<0.1

a. All variables were standardized and normalized so that an increase in investments/monitoring was positive.

b. Average across the following investments: instructing the child to work on their homework, helping the child with their homework, asking others to help the child with homework, giving the child a light source to study at night, monitoring the child's exercise books, sending the child to school with food or water, pushing the child to attend school regularly.

c. Average across 2 chores measures: hours of chores and # times fetched water.

Table 7: Treatment effect heterogeneity by parent education: Longer-run outcomes

	Heterogeneity by parent education in:							
	A. Treatment effect heterogeneity by performance − linear spec				B. Ave. treatment effect			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Independent Variable:	Treat	Treat × Score	Treat × Parent Yrs. Educ.	Treat × Score × Parent Yrs. Educ.	Treat	Treat × Parent Yrs. Educ.	Control group mean: Below-median parent educ.	Control group mean: Above-median parent educ.
Dependent Variables								
Panel A. Dropout and transfer (from endline survey data)								
Dropout	0.068	-0.0015	-0.003	0.0001	0.002	0.0007	0.033	0.005
	[0.039]*	[0.0007]**	[0.005]	[0.0001]	[0.013]	[0.0016]		
Transfer	-0.08	0.002	0.024	-0.0005	0.02	0.002	0.056	0.059
	[0.065]	[0.0013]*	[0.013]*	[0.0003]*	[0.025]	[0.005]		
Panel B. Investments (from endline survey data)								
Total educ. expenditures	1,027	-31.04	-185.69	6.219	-448.76	122.38	2,089	2,653
	[562.97]*	[14.21]**	[125.02]	[3.471]*	[347.71]	[91.91]		
ln(Total educ. expenditures)	0.369	-0.009	-0.056	0.0014	-0.049	0.012	7.293	7.489
	[0.203]*	[0.004]**	[0.037]	[0.0007]**	[0.097]	[0.018]		
Avg. std. effect across other non-monetary investments[a,b]	0.06	-0.0003	0.004	0	0.047	0.003	-0.091	0.075
	[0.097]	[0.0018]	[0.02]	[0.0004]	[0.045]	[0.009]		
Avg. std. effect across other chores[c]	0.124	-0.0017	-0.024	0.0006	0.037	0.005	-0.021	0.026
	[0.17]	[0.003]	[0.033]	[0.0006]	[0.073]	[0.014]		
Panel C. Attendance and grades (from data collected from schools)								
Attendance rate following baseline survey	0.08	-0.0015	-0.018	0.0003	0.011	-0.003	0.894	0.927
	[0.044]*	[0.0008]*	[0.008]**	[0.0001]**	[0.014]	[0.003]		
End-of-year grade	-0.0013	0.0003	0.03	-0.0008	0.008	-0.005	1.94	1.99
	[0.161]	[0.004]	[0.031]	[0.0007]	[0.059]	[0.012]		

Notes. Data sources are endline survey and endline data collected from schools. Each observation is a child. Standard errors clustered at the household level. All regressions control for child baseline achievement, grade fixed effects, school fixed effects, parents' education, the between-sibling achievement gap, and the baseline value of the dependent variable, if available (not available for dropouts, transfers, pushing children to attend school). Parent Yrs. Educ. is average years of education across parents in the household. *** p<0.01, ** p<0.05, * p<0.1
a. All variables were standardized and normalized so that an increase in investments/monitoring was positive.
b. Average across the following investments: instructing the child to work on their homework, helping the child with their homework, asking others to help the child with homework, giving the child a light source to study at night, monitoring the child's exercise books, sending the child to school with food or water, pushing the child to attend school regularly.
c. Average across 2 chores measures: hours of chores and # times fetched water.

Appendix Figure 1: Inaccurate Beliefs about Children's Relative Achievement

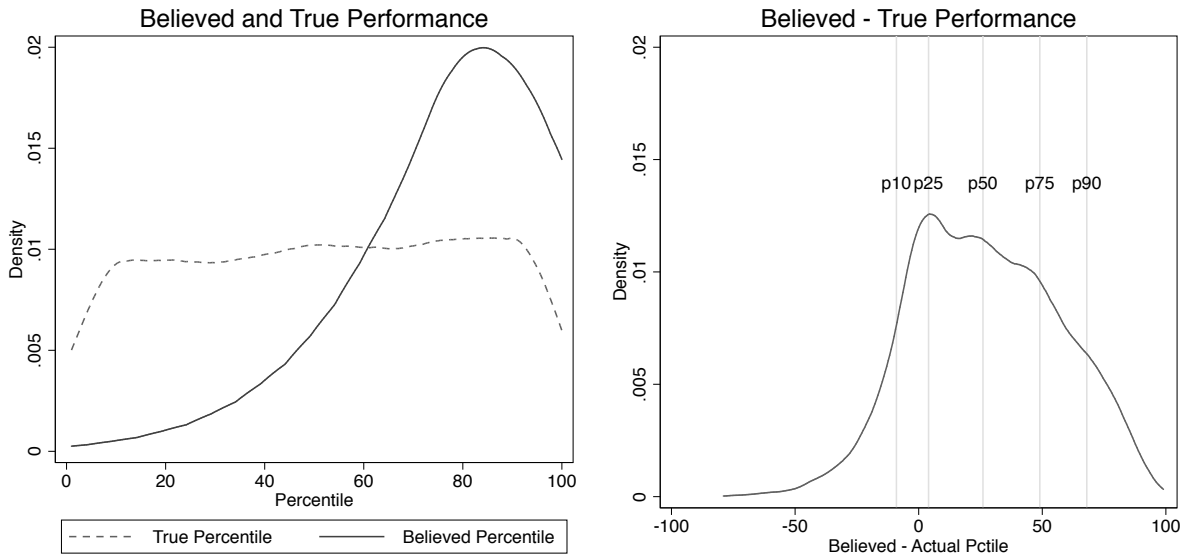

Notes: Data source is baseline data (full sample). The left graph shows kernel density plots comparing the distribution of parents' beliefs about their children's Term 2 2011-2012 relative achievement (i.e., within-class percentile rank), elicited at the beginning of the baseline survey, with the distribution of their children's true Term 2 relative achievement. The right graph shows a kernel density plot of the distribution, across parents, of each parent's beliefs about their child's relative achievement minus their child's true relative achievement. The lines represent the percentiles of the distribution.

Appendix Figure 2: Lottery Ticket Allocations

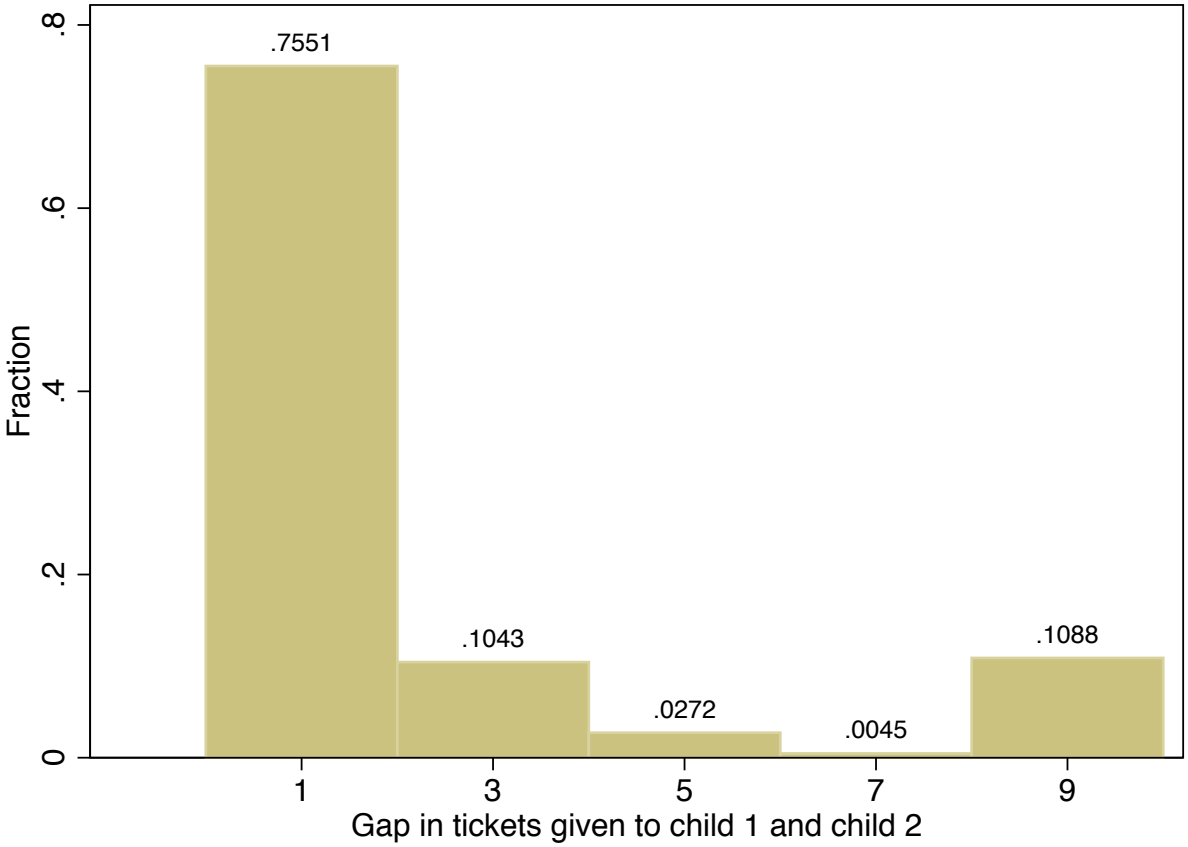

Notes: Data source is baseline data (full sample). Histogram shows how the parents split their lottery tickets between their children and, specifically, the number of tickets given to the child who received more tickets relative to the number of tickets given to the child who received fewer tickets. The total number of tickets was 9.

Appendix Table 1: Parents' behavior responds more to positive belief shocks than negative belief shocks

	Math Workbook		**English Workbook**		**Updating**	
Dependent Variable:	100 for Advanced, 0 for Average, -100 for Beginner		100 for Advanced, 0 for Average, -100 for Beginner		Post-intervention beliefs	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × (True − Believed Score)×(Pos. Shock)	0.742**	0.909*	0.650**	0.561	0.187	0.420**
	[0.355]	[0.519]	[0.278]	[0.706]	[0.116]	[0.167]
Treat × (True − Believed Score)	1.246***	1.054***	1.342***	1.302***	0.462***	0.470***
	[0.122]	[0.243]	[0.142]	[0.275]	[0.042]	[0.108]
Score Used	Math	Math	English	English	Overall	Overall
Sample: People with intermediate baseline beliefs only		✓		✓		✓
Observations	5,233	1,732	5,233	1,679	5,240	1,330
R-squared	0.136	0.179	0.144	0.172	0.170	0.249

Notes. Standard errors in brackets. Standard errors clustered at household level. Regressions control for school FE, parents' education, child achievement, the between-child score gap, and all of the main and interaction terms (i.e., Treat, (True - Believed Score), Positive shock, and all of their double and triple interactions). Columns (2), (4), and (6) restrict the sample to parents whose predicted behavior based on baseline beliefs would be near the middle of the range of potential outcome variables (i.e., parents who, based on their baseline beliefs, would be predicted to choose the average workbooks, or who have baseline beliefs near the median).
*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 2: Heterogeneity in treatment effects by achievement and beliefs

	Math Workbook	English Workbook	Textbook WTP	Lottery	
Dependent Variable:	100 for Advanced, 0 for Average, -100 for Beginner		ln(Textbook WTP) for Math−English	Tickets	
	(1)	(2)	(3)	(4)	(5)
Treat × True score	1.64***	1.67***	-0.015***	0.048***	
	[0.086]	[0.086]	[0.0020]	[0.0061]	
Treat × Believed score	-1.52***	-1.55***	0.011***	-0.034***	
	[0.094]	[0.082]	[0.0018]	[0.0067]	
True score	0.16***	0.13**	-0.000012	0.0019	
	[0.057]	[0.059]	[0.0015]	[0.0041]	
Believed score	2.22***	2.33***	-0.021***	0.070***	
	[0.065]	[0.055]	[0.0012]	[0.0045]	
Treat	-5.35	-0.11	0.046		
	[6.35]	[4.85]	[0.041]		
Treat × Higher-scoring Sibling					1.27***
					[0.14]
Treat × Believed Higher-scoring Sib					-0.83***
					[0.14]
Higher-scoring Sibling					-0.016
					[0.091]
Believed Higher-scoring Sibling					1.55***
					[0.091]
Household FE					✓
Observations	5,233	5,233	5,177	5,214	5,212
R-squared	0.374	0.405	0.097	0.209	0.212
P-val: (Treat × True) + (Treat × Beliefs)=0	0.231	0.197	0.098	0.024	
P-val: (Treat × High-Perf Sib) + (Treat × Bel'v'd High-Perf Sib)=0					0.002

Notes. Robust standard errors in brackets. Standard errors clustered at the household level. Regressions control for school FE, parents' education, child achievement, and the between-child score gap.
*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 3: Detailed skills treatment and uncertainty effects

	Coefficient estimate (std. error) for:				
	A. Ave. Treatment effects	B. Detailed skills treatment effects		C. Uncertainty: Beliefs within 10 pts of truth)	
Independent Variable:	Treat	Treat	Detailed Skills Treatment	Treat	N
Dependent Variables					
Panel A. Experimental outcomes					
Math worbook	-31.28	-30.44	-1.657	-6.56	5,239
	[2.077]***	[2.595]***	[3]	[3.638]*	
English workbook	-12.55	-11.24	-2.572	2.328	5,239
	[2.153]***	[2.587]***	[2.886]	[3.205]	
ln(WTP for Math - English textbook)	0.14	0.213	-0.143	0.07	5,219
	[0.041]***	[0.049]***	[0.057]	[0.072]	
Lottery tix given: higher - lower performer	0.984	0.949	0.069	0.478	2,629
	[0.128]***	[0.154]***	[0.178]	[0.192]***	
Panel B. Longer-run data: Dropout and transfer					
Dropout	0.004	-0.004	0.016	0.004	1,786
	[0.007]	[0.008]	[0.01]	[0.009]	
Transfer	0.03	0.045	-0.029	0.029	1,781
	[0.014]**	[0.018]***	[0.02]	[0.024]	
Panel C. Longer-run data: Investments					
Total educ. expenditures	104.45	2.632	192.26	378.10	1,729
	[164.32]	[188.74]	[260.81]	[274.06]	
ln(Total educ. expenditures)	0.0013	0.031	-0.056	0.07	1,709
	[0.049]	[0.06]	[0.071]	[0.086]	
Avg. std. effect: non-monetary investments[a,b]	0.065	0.045	0.037	0.075	1,720
	[0.026]***	[0.032]	[0.035]	[0.04]*	
Avg. std. effect: other chores[c]	0.058	0.037	0.04	0.051	1,681
	[0.041]	[0.051]	[0.059]	[0.07]	
Panel D. Longer-run data: Attendance and grades					
Attendance rate	-0.002	-0.006	0.007	0.01	1,827
	[0.008]	[0.009]	[0.011]	[0.013]	
End-of-year grade	-0.016	-0.026	0.02	0.019	1,241
	[0.036]	[0.047]	[0.052]	[0.068]	

Notes. Data sources are baseline survey, endline survey and endline data collected from schools. Standard errors clustered at the household level. Regressions control for school FE, parents' education, child achievement, the between-child score gap, and the baseline value of the dependent variable (not available for dropouts, transfers, pushing children to attend school). *** p<0.01, ** p<0.05, * p<0.1

a - c. See notes for Tbl 7.

Appendix Table 4: The finding that less-educated parents have less accurate beliefs is robust to different measures of education and achievement

		Coefficient estimate for:							
		Respondent's				Parent Average			
	Full Sample Mean [SD]	Years of educ.	Above-median educ.	At least sec-ondary educ.	Parent is literate	Years of educ.	Above-median educ.	At least sec-ondary educ.	Parent is literate
Dependent Variables									
Panel A. Scores									
Abs Val [Believed − True Overall Score]	20.39	-0.18	-0.831	-2.549	-1.082	-0.197	-0.488	-3.526	-0.996
	[14.47]	[0.063]***	[0.451]*	[0.684]***	[0.491]**	[0.066]***	[0.45]	[0.72]***	[0.593]*
Abs Val [Believed − True Math Score]	25.75	-0.246	-0.816	-4.056	-0.818	-0.274	-0.733	-4.954	-1.192
	[18]	[0.073]***	[0.538]	[0.787]***	[0.582]	[0.079]***	[0.537]	[0.867]***	[0.713]*
Abs Val [Believed − True English Score]	21.35	-0.106	-0.781	-1.328	-0.898	-0.105	-0.326	-1.869	-0.846
	[16.44]	[0.071]	[0.494]	[0.767]*	[0.548]	[0.074]	[0.493]	[0.832]**	[0.649]
Abs Val [Believed − True Chichewa Score]	23.81	-0.248	-1.458	-2.921	-0.725	-0.299	-1.513	-4.204	-0.517
	[17.54]	[0.075]***	[0.537]***	[0.848]***	[0.571]	[0.077]***	[0.534]***	[0.864]***	[0.669]
Abs Val [Believed − True (Math-English) Score]	22.08	-0.061	-0.134	-1.735	0.723	-0.035	-0.482	-1.461	0.574
	[17.4]	[0.072]	[0.526]	[0.754]**	[0.58]	[0.078]	[0.525]	[0.834]*	[0.698]
Abs Val [Believed − True Overall Score (Child1-2)]	18.67	-0.204	-1.944	-0.544	-1.307	-0.244	-1.568	-1.768	-1.717
	[15.13]	[0.084]***	[0.6]***	[0.936]	[0.67]*	[0.091]***	[0.599]***	[0.996]*	[0.814]**
Wrong about which child higher scoring	0.311	-0.007	-0.057	-0.026	-0.016	-0.008	-0.048	-0.054	-0.022
	[0.463]	[0.003]***	[0.018]***	[0.03]	[0.02]	[0.003]***	[0.018]***	[0.032]*	[0.024]
Panel B. Percentiles									
Abs Val [Believed − True Overall Percentile]	32.16	-0.355	-1.99	-4.9	-2.654	-0.396	-1.61	-5.873	-2.78
	[24.03]	[0.098]***	[0.704]***	[1.113]***	[0.749]***	[0.105]***	[0.701]**	[1.155]***	[0.933]***
Abs Val [Believed − True Math Percentile]	33.34	-0.372	-1.928	-5.82	-2.671	-0.413	-1.885	-6.861	-2.848
	[25]	[0.101]***	[0.73]***	[1.109]***	[0.801]***	[0.11]***	[0.726]***	[1.187]***	[0.992]***
Abs Val [Believed − True English Percentile]	30.58	-0.233	-1.514	-2.377	-2.147	-0.292	-1.176	-3.354	-2.493
	[23.35]	[0.097]***	[0.687]**	[1.139]**	[0.73]***	[0.105]***	[0.682]*	[1.221]***	[0.92]***
Abs Val [Believed − True Chichewa Percentile]	33.78	-0.251	-1.014	-3.926	-1.411	-0.293	-0.952	-5.032	-1.512
	[24.72]	[0.101]***	[0.728]	[1.143]***	[0.775]*	[0.109]***	[0.724]	[1.217]***	[0.943]
Abs Val [Believed − True (Math-English) Percentile]	25.66	-0.314	-2.373	-2.059	-1.156	-0.287	-2.183	-2.2	-1.428
	[21.56]	[0.09]***	[0.638]***	[1.002]**	[0.701]*	[0.096]***	[0.637]***	[1.078]**	[0.848]*
Abs Val [Believed − True Overall Percentile (Child1-2)]	32.55	-0.448	-3.525	-3.766	-3.44	-0.473	-2.185	-4.948	-3.234
	[22.74]	[0.125]***	[0.904]***	[1.353]***	[0.99]***	[0.133]***	[0.902]***	[1.444]***	[1.2]***
Wrong about which child higher percentile	0.339	-0.007	-0.064	-0.051	-0.036	-0.008	-0.051	-0.075	-0.027
	[0.473]	[0.003]***	[0.019]***	[0.03]*	[0.02]*	[0.003]***	[0.019]***	[0.032]**	[0.025]
Sample size	5,268	5,230	5,230	5,230	5,230	5,230	5,230	5,230	5,242

Notes. Each observation is a child. Standard errors clustered at the household level. Regressions control for child's gender, grade, parent gender.
*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 5: Heterogeneity in belief accuracy by parent education: Variation within schools and conditional on achievement

	Dependent Variable= Abs Val [True − Believed Achievement]					
	(1)	(2)	(3)	(4)	(5)	(6)
Performance Measure:	Overall	Overall	Math	Math	English	English
Parents' years of education	-0.096	-0.072	-0.285***	-0.254***	-0.048	-0.014
	[0.069]	[0.059]	[0.084]	[0.073]	[0.079]	[0.072]
Observations	5,019	5,019	5,021	5,021	5,021	5,021
Dep var mean	20.41		25.82		21.42	
	(7)	(8)	(9)	(10)	(11)	(12)
Performance Measure:	Chichewa	Chichewa	Math−Eng	Math−Eng	Child 2−1	Child 2−1
Parents' years of education	-0.165*	-0.126*	-0.086	-0.155**	-0.162	-0.239***
	[0.084]	[0.068]	[0.085]	[0.076]	[0.102]	[0.091]
Observations	5,021	5,021	5,021	5,021	2,514	2,514
Dep var mean		22.12		18.68		
Col. Specification Details						
School FE	✓		✓		✓	
Achievement Control		✓		✓		✓

Notes. Robust standard errors in brackets. Standard errors clustered at the household level. All columns include controls for child gender, grade FE, parent gender, and whether the parent is the primary education decisionmaker. Parent years of education is average across parents in the household. *** p<0.01, ** p<0.05, * p<0.1

Appendix Table 6: Heterogeneity by parent education in overconfidence, uncertainty, and achievement

	Overconfidence			**Uncertainty**			**Achievement**	
Dependent Variable:	Believed - True Score			Std. Dev. of Beliefs about Score			Score	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Parent's Years of education	-0.079	-0.078	0.156**	-0.614***	-0.611***	-0.609***	0.348***	0.354***
	[0.088]	[0.090]	[0.071]	[0.055]	[0.056]	[0.055]	[0.076]	[0.077]
Child and parent controls		✓	✓		✓	✓		✓
Score control			✓			✓		
Observations	5,220	5,019	5,220	5,171	4,974	5,171	5,230	5,029
R-squared	0.000	0.002	0.368	0.039	0.042	0.040	0.004	0.007
Dep. Var. Mean	15.63			7.66			46.72	

Notes. Standard errors in brackets. Standard errors clustered at household level.
*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 7: Effect of Information on Endline Outcomes (Detailed data)

	A. Heterogeneity in treatment effects by performance – linear spec		B. Heterogeneity in treatment effects by performance – nonparametric spec		C. Avg. treatment effect		
Independent Variable:	Treat	Treat	Treat× Score	Treat	Treat× Above-Median Score	Control group mean	N
Dependent Variables							
Panel A. Dropout and transfers (from endline survey data)							
Dropout	0.055	-0.0011	0.022	-0.037	0.004	0.021	1,786
	[0.021]***	[0.0004]***	[0.012]*	[0.015]***	[0.007]		
Transfer	0.023	0.0002	0.022	0.017	0.03	0.057	1,781
	[0.036]	[0.0007]	[0.019]	[0.025]	[0.014]**		
Panel B. Education Expenditures (from endline survey data)							
Total educ. expenditures	119.70	-0.325	100.54	4.181	104.45	2,362.06	1,729
	[291.50]	[6.841]	[177.56]	[230.00]	[164.32]		
ln(Total educ. expenditures)	0.093	-0.0019	0.014	-0.03	0.0013	7.389	1,709
	[0.114]	[0.002]	[0.061]	[0.074]	[0.049]		
Expenditures on school fees	125.67	-2.924	-3.618	-16.15	-11.26	452.53	1,729
	[63.26]**	[1.612]*	[27.18]	[47.65]	[30.82]		
Supplementary educ. expenditures	-67.75	3.623	79.80	43.81	101.94	1,902.92	1,729
	[275.83]	[6.44]	[172.71]	[217.88]	[157.24]		
Books and school supplies	105.10	-0.959	85.90	-52.65	60.20	617.64	1,729
	[94.08]	[1.803]	[59.99]	[63.27]	[57.42]		
Uniforms	-93.89	2.737	-25.07	118.63	34.34	806.40	1,729
	[138.75]	[2.5]	[87.79]	[100.21]	[70.23]		
Backpacks	7.431	0.691	43.20	-6.653	39.77	178.61	1,729
	[50.97]	[1.099]	[31.53]	[36.16]	[27.08]		
Tutoring	0.771	-0.794	-1.68	-71.15	-36.40	300.27	1,729
	[158.37]	[4.308]	[83.13]	[151.06]	[88.82]		
Panel C. Non-monetary investments (from endline survey data)							
Helped child with homework	-0.034	0.0001	-0.046	0.033	-0.03	0.374	1,699
	[0.063]	[0.0012]	[0.034]	[0.042]	[0.028]		
Asked someone to help child with homework	0.1	-0.001	0.07	-0.033	0.055	0.243	1,710
	[0.064]	[0.0013]	[0.033]**	[0.042]	[0.027]**		
# times gave child light source to study at night over last 4 weeks	0.207	0.005	0.274	0.316	0.425	2.61	1,674
	[0.889]	[0.018]	[0.481]	[0.616]	[0.402]		
# times child went to school without food or water in last 4 weeks	-2.374	0.019	-1.778	0.654	-1.461	10.68	1,677
	[1.198]**	[0.022]	[0.671]***	[0.772]	[0.543]***		
Has to push child to attend school regularly	0.028	0.0008	0.059	0.017	0.067	0.341	1,666
	[0.062]	[0.0012]	[0.034]*	[0.041]	[0.026]***		
# times monitored child's exercise books in last 4 weeks	-1.352	0.005	-1.12	0.002	-1.132	8.458	1,681
	[1.133]	[0.022]	[0.613]*	[0.735]	[0.486]**		
# times instructed child to work on homework in last 4 weeks	0.819	-0.005	0.447	0.224	0.559	1.972	1,669
	[0.472]*	[0.009]	[0.29]	[0.345]	[0.249]**		
Ave. std. effect across other investments[b]	0.07	-0.0001	0.057	0.015	0.065	-0.012	1,720
	[0.057]	[0.0011]	[0.032]*	[0.039]	[0.026]***		
Panel D. Chores (from endline survey data)							
Hours of chores given to child over last 4 weeks	0.546	0.029	1.43	1.008	1.905	23.81	1,676
	[2.936]	[0.066]	[1.385]	[2.182]	[1.325]		
# times child fetched water in last 4 weeks	0.155	0.003	0.061	0.42	0.273	4.656	1,671
	[1.016]	[0.02]	[0.512]	[0.619]	[0.37]		
Ave. std. effect across chores[c]	0.01	0.001	0.034	0.049	0.058	-0.0009	1,681
	[0.104]	[0.002]	[0.05]	[0.069]	[0.041]		
Panel E. Attendance and grades (from data collected from schools)							
Attendance rate following baseline survey	-0.008	0.0001	-0.0015	-0.0017	-0.002	0.911	1,827
	[0.026]	[0.0005]	[0.012]	[0.015]	[0.008]		
End-of-year grade	0.122	-0.003	0.03	-0.095	-0.016	1.97	1,241
	[0.091]	[0.0019]	[0.047]	[0.07]	[0.036]		

Notes. Data sources are endline survey and endline data collected from schools. Each observation is a child. Standard errors clustered at the household level. Regressions control for school FE, parents' education, child achievement, the between-child score gap, and the baseline value of the dependent variable (not available for dropouts, transfers, pushing children to attend school). *** p<0.01, ** p<0.05, * p<0.1

a. All variables were standardized and normalized so that an increase in investments/monitoring was positive.

b. Average across the following investments: instructing the child to work on their homework, helping the child with their homework, asking others to help the child with homework, giving the child a light source to study at night, monitoring the child's exercise books, sending the child to school with food or water, pushing the child to attend school regularly.

c. Average across 2 chores measures: hours of chores and # times fetched water.

A Sample Information Intervention Report Card

Report Card

Name: **NDEMA LONGWE** **Standard:** **2**

	Score	Grade	Position
Maths:	75/100	3	10/100
English:	33/100	1	71/100
Chichewa:	67/100	3	38/100
Overall:	58/100	2	52/100

Number of Exams Administered in Class: 3

Grades
1 = Needs support
2 = Average
3 = Good
4 = Excellent