

Parents' Beliefs about Their Children's Academic Ability: Implications for Educational Investments[†]

By REBECCA DIZON-ROSS*

Schools worldwide distribute information to parents about their children's academic performance. Do frictions prevent parents, particularly low-income parents, from accessing this information to make decisions? A field experiment in Malawi shows that, at baseline, parents' beliefs about their children's academic performance are often inaccurate. Providing parents with clear, digestible performance information causes them to update their beliefs and adjust their investments: they increase the school enrollment of their higher-performing children, decrease the enrollment of lower-performing children, and choose educational inputs that are more closely matched to their children's academic level. Heterogeneity analysis suggests information frictions are worse among the poor. (JEL C93, D83, I21, I24, J13, O15)

It is commonly believed that one reason poor households remain poor is that they lack information (World Bank 1998, 2008). Indeed, there is extensive evidence that providing information to poor individuals affects their decision-making across many domains including health (e.g., Dupas 2011), education (e.g., Jensen 2010), and labor supply (e.g., Chetty and Saez 2009). However, the vast majority of this evidence concerns information that even richer households may not have or use, such as the economic returns to education or market-level summary statistics. The question remains whether frictions prevent poor individuals from taking advantage

*Booth School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637 (email: rrdr@chicagobooth.edu). Esther Duflo was the coeditor for this article. I am very grateful to Pascaline Dupas, Caroline Hoxby, and Seema Jayachandran for their guidance, and to Ran Abramitzky, Abhijit Banerjee, Jim Berry, Marianne Bertrand, Nick Bloom, Doug Bernheim, Eric Budish, Manasi Deshpande, Celine Dizon, Elise Dizon-Ross, Natalie Douvos, Alex Eble, Liran Einav, Nick Hagerty, Rema Hanna, Johannes Haushofer, Yael Hochberg, Anil Jain, Asim Khwaja, Anjini Kochar, Dan Lee, Shirlee Lichtman, Matthew Lowe, Rachael Meager, Ben Olken, Arianna Ornaghi, Jonah Rockoff, Sheldon Ross, Ashish Shenoy, Fabiana Silva, Melanie Wasserman, Tom Wollmann, Jenny Ying, Owen Zidar, and workshop participants at Stanford, MIT, and Harvard, and seminar participants at Princeton, Yale, Columbia, University of Chicago, Chicago Booth, Northwestern, Harvard Business School, Stanford GSB, UCLA, UCSD, Stanford, Columbia Teacher's College, Duke, NBER Development Fall 2014, NBER Education Spring 2014, PacDEV 2014, and NEUDC 2013 for helpful comments and discussions. I thank Bridget Hoffmann, Rachel Levenson, and Michael Roscitt for help with the fieldwork, and Christine Cai and Yashna Nandan for excellent research assistance. 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, SIEPR, the Stanford Economics Department, and the DDRO and GRO Funds. This study is registered in the AEA RCT Registry under the unique identifying number AEARCTR-0001808. I declare that I have no relevant or material financial interests that relate to the research described in this paper. All errors are my own.

[†]Go to <https://doi.org/10.1257/aer.20171172> to visit the article page for additional materials and author disclosure statement.

of information that seems available to all, but which perhaps only those in richer households or countries can access and exploit.

I examine these issues in a high-stakes environment: parents making decisions about their children's education. It is widely presumed that correct educational decisions (such as whether to go to college or sign up for remedial tutoring) vary across individuals. Anecdotally, in developing countries, the most important child-specific factor determining parents' educational decisions is their children's school performance. School performance information appears to be freely available: schools worldwide deliver report cards to parents that contain this information, and parents can also observe their children's academic skills directly. And yet, there may be barriers preventing some parents from making use of this information (Banerjee et al. 2010). For example, parents in developing countries are often illiterate and may not be able to read or understand report cards. Limited education may also make it difficult for parents to judge their children's performance themselves, especially if their children go further in school than they did, as is common in developing countries.¹

This paper establishes that there are in fact substantial and consequential information frictions among poor parents regarding their children's school performance.² First, I show that many parents in a developing country context have inaccurate beliefs about their children's school performance. Second, I demonstrate that parents base important educational decisions upon their inaccurate beliefs even though they would prefer to use the correct information that is, in principle, readily available. Third, I use a randomized experiment to show that a simple informational intervention can significantly alleviate the impacts of limited information: providing information directly to parents in a clear and digestible way causes them to update their beliefs and adjust their decisions accordingly. Finally, I provide evidence on a link between information barriers and poverty, showing that poorer, less-educated parents have less accurate baseline beliefs than richer parents, and that their beliefs and certain of their investments respond more to information.

I demonstrate these findings by conducting a randomized field experiment in Malawi. The experiment delivers information to randomly-selected parents with children in primary school about their children's "academic performance," which hereafter refers to average performance on achievement tests administered by schools during the term before the intervention; on average, schools offered four tests per subject in three subjects.³ This information is delivered verbally and in a clear manner. I measure the effect of the information on parents' beliefs and on a broad range of their investments and decisions, including both a series of real-stakes investment options and decisions presented to parents through the experiment ("experimental outcomes"), and more traditional endline outcomes such as enrollment and attendance in school ("non-experimental outcomes"). The analysis proceeds as follows.

¹ Free primary schooling in most developing countries only became widely available in the last 10–20 years, and the average adult in sub-Saharan Africa has fewer than five years of education (UNESCO 2013).

² Although this paper does not claim to distinguish between the various frictions at play, they may include parental illiteracy, limited parental ability to directly assess academic skills, the complexity of existing report cards, and problems with the report card delivery mechanism.

³ As described in Section II, all schools send report cards to parents with this information; thus, academic performance information should in theory be freely available in this context.

I use baseline beliefs data to establish my first finding: that parents' beliefs are inaccurate. On average, parents' beliefs about academic performance diverge from true performance by more than one standard deviation of the performance distribution. When comparing two of their children, one-third of parents are mistaken about which child is higher-performing.

Next, I combine information on believed performance, true performance, and investment decisions to test whether inaccurate beliefs affect parents' decisions. I establish my second finding—that at baseline, parents base important decisions on their inaccurate beliefs—by demonstrating that in the control group, the relationship between believed performance and investments is stronger than the relationship between true performance and investments. I then establish my third finding—that the information intervention reduces knowledge barriers—by showing that in the treatment group, the relationship between true performance and investments becomes stronger, increasing to resemble the relationship in the control group between *believed* performance and investment. This is because parents' beliefs become more closely aligned with true performance, and they adjust their investments accordingly. The analyses show that student performance is an important input into parents' decisions, but that parents are often quite wrong about performance, resulting in important investment “mistakes” (i.e., wedges between how parents would like to allocate their investments given their children's true academic performance and how they allocate them in reality).

I establish two broad categories of investment mistakes. The first is misallocation in the *level* of investment across children, i.e., cases in which the total amount invested in each child is not what parents would want given children's academic performance. I test for this type of mistake using enrollment in primary school and an experimental outcome proxying for resources allocated toward secondary school. Providing information has impacts on both, causing reallocations toward higher-performing students. These results suggest that parents prefer to allocate more years of schooling to their higher performers. The analysis clearly shows that information frictions affect investments; the implications for welfare then depend on whether there are other interacting market frictions, including whether parents are correct about the education production function. I discuss this issue later in the paper.

The second category of mistake I uncover is misallocation in the *types* of investment chosen for a given child, i.e., failures to tailor the specific input mix correctly to a child's academic level, such as purchasing an advanced textbook for a low-performing child when the remedial textbook would have benefited her more. Here I use several experimental outcomes, such as the demand for books designed for students of different performance levels. The prediction is that returns will be higher if the level of the selected book matches the level of the child's performance, and I find that providing information triples the closeness of the match. These types of parental decisions are now more relevant than ever in developing countries, since the use of supplementary inputs is growing rapidly (Paviot, Heinsohn, and Korkman 2008).

Finally, I provide evidence on a link between poverty and information barriers, showing that poorer, less-educated parents have less accurate baseline beliefs than richer parents, and that they adjust their beliefs and certain investments more in response to information.

This paper contributes to our understanding of how information frictions affect decisions. A large literature shows that providing information to households affects decision-making across many domains.⁴ However these interventions have primarily delivered information that one might not expect households to know, even richer or well-informed households. Some papers provide information that is difficult or sometimes impossible to obtain, such as (normally unobservable) economic returns to an activity (e.g., Jensen 2010), or statistics that require another entity's private information, such as school-average student performance or normally-unreleased data on student effort in school (Andrabi, Das, and Khwaja 2016; Bergman 2016). Others deliver information that is technically available but requires nontrivial effort to obtain, such as personalized information about the cost of health plans or the Earned Income Tax Credit (EITC) schedule (Chetty and Saez 2009, Kling et al. 2011). In contrast, this paper examines information that is seemingly readily available and shows that it is still not incorporated into the decisions made by the poor. This is important not just because it demonstrates a more surprising failure of optimization, but also because it provides evidence of a channel through which the poor might remain poor: frictions that prevent them from using information that seems readily available to all, but that only the rich can leverage.

This paper also contributes to the literature on information frictions in education. This literature has focused on misinformation about aggregate factors, such as the population-average returns to education, school quality, or other features of the education system,⁵ thus abstracting away from the fact that correct individual education decisions (such as whether to go to college or whether to invest in a remedial textbook) vary across individuals. Here, I shift focus from aggregates to the heterogeneity within the population, providing, to my knowledge, the first evidence using exogenous variation in beliefs to establish a causal link between misinformation about individual-level characteristics and investment decisions.⁶

The paper proceeds as follows. Section I motivates the empirical approach. Section II describes the context and experimental design. Section III presents the results, and Section IV concludes.

I. Empirical Approach for Detecting Mistakes

In this section, I present a simple framework in which inaccurate beliefs cause mistakes and use it to generate empirical predictions. I then discuss how to use an experiment to test the predictions. The framework moves beyond providing information and examining the average treatment effect (ATE) on investments, as ATEs can produce a biased picture of the effects of misinformation. For example, if providing information affects the salience of education, it might generate a nonzero

⁴For example, Liebman and Luttmer (2015); Kling et al. (2011); Bhargava and Manoli (2015); Duflo and Saez (2002); Dupas (2011); Fitzsimons et al. (2016); Jamison, Karlan, and Rafter (2013); Beshears et al. (2015).

⁵See, for example, Jensen (2010); Nguyen (2008); Andrabi, Das, and Khwaja (2016); Bettinger et al. (2012); Dinkelman and Martínez A (2014); Hoxby and Turner (2013); Wiswall and Zafar (2015).

⁶This builds on prior studies that use observational data to show that students' beliefs about their own abilities predict their decisions, such as college major choice or college dropout (Chevalier et al. 2009; Arcidiacono, Hotz, and Kang 2012; Stinebrickner and Stinebrickner 2012, 2014). My findings also complement a recent information experiment by Bobba and Frisncho (2016) that tests predictions about the differential roles of the mean and variance of beliefs on educational decisions.

ATE even if baseline information was perfect. ATEs can also understate information distortions: only beliefs that are biased at the population level will produce a non-zero ATE, while beliefs that are individually inaccurate but not biased on average can still produce important mistakes.⁷

Consider a parent choosing investments in her children's schooling. Loosely speaking, she chooses both the *level* of spending on each child and the specific *type* of educational resources for each child, for example, what difficulty level of textbook or tutoring to choose. Denote one of parent i 's investment choices for child j as s_{ij} . The perceived production function for child j 's "quality" (i.e., human capital or expected lifetime earnings) is

$$(1) \quad q_{ij} = f(s_{ij}, a_{ij})$$

with a_{ij} denoting child j 's baseline academic performance and f concave in s_{ij} .⁸ A key assumption, which can later be tested in the data, is that parents perceive the returns to input s_{ij} to vary with a_{ij} ,

$$(2) \quad \frac{\partial^2 f}{\partial s \partial a} \neq 0.$$

If parent i has perfect information about a_{ij} , she chooses s_{ij} to maximize household utility:

$$(3) \quad s_{ij}^* = \arg\max_s U(q_{ij})$$

subject to a budget constraint. Given equation (2), s_{ij}^* depends on a_{ij} . I can thus define the *preferred investment function*, $s^*(a)$, as the full set of solutions to equation (3) for all values of a in the population. Much of the analysis centers around the derivative of this function ($\partial s^* / \partial a$), which may vary across investments. For example, if parents' utility functions maximize returns and s is a perceived substitute with performance ($\partial^2 f / \partial s \partial a < 0$), then $\partial s^* / \partial a < 0$; if s is a perceived complement with performance, $\partial s^* / \partial a > 0$.⁹ I discuss the predictions for $\partial s^* / \partial a$ for each investment as I proceed through the analysis.

Now, assume parent i does not know child j 's true performance a_{ij} . Instead, her beliefs about a_{ij} are described by the distribution $g(\alpha_{ij}, \sigma_{ij}^2)$, with α_{ij} her mean beliefs and σ_{ij}^2 her belief uncertainty. She thus chooses s_{ij} to maximize expected utility taken over $g(\alpha_{ij}, \sigma_{ij}^2)$. Under some models, only α_{ij} (and not σ_{ij}^2) would affect her choice of s_{ij} . For example, σ_{ij}^2 would not matter if parents' maximand takes a quadratic loss

⁷ Even if parents' beliefs about their children's *overall* performance are biased, their beliefs about the performance measure relevant for a given decision may not be. For example, if a parent can only afford to send one child to school and wants to send her higher-performer, the relevant metric is her belief about her children's performance *relative* to each other, which by definition is mean-zero in any household and so cannot be biased on average, even if parents are overconfident.

⁸ Note that we take a_{ij} to be a baseline measure which affects the returns to the investment; it is an input to the efficacy of s_{ij} , not an outcome affected by s_{ij} .

⁹ Maximizing returns corresponds to the case where the utility function is linear in different children's qualities. If parents care not just about maximizing returns but also about equalizing between their children, then $\partial s^* / \partial a$ will also depend on the utility function cross-partials in the different children's qualities.

form: $U(f(s_{ij}, a_{ij})) = -c(\gamma s_{ij} - a_{ij})^2$. For expositional simplicity, I first restrict attention to this case before generalizing below. Here, parent i 's chosen investment, given her beliefs $g(\alpha_{ij}, \sigma_{ij}^2)$, equals $s^*(\alpha_{ij})$; if mean beliefs are inaccurate ($\alpha_{ij} \neq a_{ij}$), this choice, $s^*(\alpha_{ij})$, diverges from the utility-maximizing choice, $s^*(a_{ij})$, a "mistake" which causes her utility to be inefficiently low.

If we have data on α , a , and s , what empirical patterns would suggest that parents are making mistakes? Mistakes happen because s does not vary with a according to the preferred function $s^*(\cdot)$. Defining the *actual investment function* $\tilde{s}(\cdot)$ as the conditional expectation, taken across individuals, of investments chosen as a function of *true* performance, $\tilde{s}(a) \equiv E(s|a)$, we want to test for a divergence between $\tilde{s}(\cdot)$ and $s^*(\cdot)$. The empirical analog of $s^*(\cdot)$ is the conditional expectation of investments, s , given α (instead of a).

The form of the divergence between $\tilde{s}(\cdot)$ and $s^*(\cdot)$ depends on the joint distribution of α and a . In most beliefs data about performance measured on a bounded scale (e.g., Banerjee et al. 2010, Alexander and Entwisle 2006), including the data used in this paper, an empirical signature of belief inaccuracies is that mean beliefs have a slope less than 1 if plotted on true performance. I call this pattern "attenuation" (panel A of Figure 1). Attenuation is a very general form of belief inaccuracy, present whenever beliefs are positively but imperfectly correlated with true performance, as long as the variance of beliefs is not too much larger than the variance of true performance. Formally, this follows from the ordinary least squares (OLS) formula:

$$(4) \quad \text{Slope of } \alpha \text{ on } a = \text{corr}(\alpha, a) \frac{\text{std}(\alpha)}{\text{std}(a)}$$

with $\text{corr}(\cdot, \cdot)$ denoting the correlation and $\text{std}(\cdot)$ the cross-person standard deviation. The lower the correlation between believed and true performance, the flatter the slope. Both mean-zero belief inaccuracies and parental overconfidence can cause attenuation; in the overconfidence case, attenuation happens because parents cannot be as overconfident at the top of the a distribution as they can at the bottom. This asymmetry can cause measurement error to be negatively correlated with a , which is a sufficient condition for attenuation.

If beliefs are inaccurate and attenuated, the following prediction allows us to test whether the inaccuracies affect investments.

PREDICTION 1: *If (i) the slope of believed performance on true performance is less than 1 and (ii) investments depend on beliefs, then the slope of investments on true performance $\partial \tilde{s} / \partial a$ will be flatter than the slope of investments on beliefs $\partial s^* / \partial \alpha$.*

See Appendix Section A for the proof and a discussion of how to test for a flatter slope when $s^*(\cdot)$ is nonlinear. The intuition is that parents choose investments based on their (inaccurate) beliefs; thus, investments are steeply sloped with α , as depicted in panel B of Figure 1 for the case where $\partial s^* / \partial \alpha > 0$. But, if we look at children who are *truly* at the top of the distribution, many of their parents underestimate their performance and so on average choose inputs appropriate for lower-performing children. Analogously, many parents of children at the bottom of the distribution




FIGURE 1. EMPIRICAL APPROACH: INACCURATE BELIEFS ABOUT PERFORMANCE CAN CAUSE THE SLOPE OF INVESTMENTS AS A FUNCTION OF PERFORMANCE TO BE FLATTER THAN THE SLOPE AS A FUNCTION OF BELIEFS

Notes: Graphs are illustrative, showing one way to test whether parents’ inaccurate beliefs affect their investments. A common type of belief inaccuracy is that beliefs will be “attenuated” on true performance, i.e., have a slope less than 1 on true performance (panel A). Parents base their investments on their potentially inaccurate beliefs, and so plotting investments on beliefs shows us parents’ “preferred” slope, i.e., the slope they would opt to choose if they knew their children’s true performance (panel B). However, because beliefs are inaccurate, and in particular, attenuated, the slope of investments as a function of children’s *true* academic performance is flatter than the slope on beliefs (panel C). The interpretation of the difference in slopes is that investments are not as well tailored to academic performance as parents would like.

choose inputs appropriate for higher-performing children. This pattern causes the slope of $\tilde{s}(\cdot)$ to be more attenuated (i.e., flatter) than the preferred slope (panel C of Figure 1) and decreases welfare.¹⁰

¹⁰ Attenuation here can be seen as analogous to attenuation bias caused by classical measurement error, but to make that analogy, one needs to think of *beliefs* as the correctly measured regressor and *true* performance as the mismeasured regressor. This is because the data generating process for s_{ij} depends on beliefs.

Estimation.—It is difficult to empirically estimate the difference between the slopes of $\tilde{s}(\cdot)$ and $s^*(\cdot)$ because neither regression line is causal. Assume that parents invest according to the model above plus an error term ε that reflects all other determinants of investment: $s_{ij} = s^*(\alpha_{ij}) + \varepsilon_{ij}$. The error term could represent a factor unrelated to academic potential, such as idiosyncratic valuation of education. It could also represent parents' beliefs about the elements of their children's academic potential not captured by recent academic performance a .¹¹ Consider comparing the slope estimated from regressing investments on α to the slope estimated from regressing investments on a . The estimated slopes could differ from the true causal slopes as a result of omitted variable bias (OVB) from the correlations between ε , and α or a . In particular, if $\alpha - a$ is correlated with ε , then the slopes of the estimated lines may differ, but only due to OVB, not due to parents making mistakes.

We can use an experiment to overcome this estimation challenge. Consider an information intervention that tells parents true performance a and changes their mean beliefs to equal a . If attenuation resulted from parents allocating s based on some measure other than a (i.e., if attenuation simply resulted from OVB and not parental mistakes), then parents' investments would not change. If instead attenuation was the result of parents' inaccurate beliefs about a causing them to make mistakes, then providing information on a would allow parents to correct their baseline mistakes and choose their preferred investment $s^*(a)$, i.e., to invest along the preferred investment function.

PREDICTION 2: *If (i) the slope of baseline investments on true performance is attenuated due to inaccurate beliefs about performance and (ii) providing information about performance causes parents to update their beliefs, then providing performance information will increase the magnitude of the slope of investments on performance.*

See Appendix Section A for the proof.

A change in the slope shows that parents think incorporating new information improves their decisions, and thus that, from their own perspectives, their decisions were not optimal to start with. Assessing the welfare implications would be more complex, relying, for example, on whether the perceived production function is correct. I discuss this in greater detail later.

Heterogeneity by Parent Socioeconomic Status (SES).—Another goal of this paper is to provide evidence on whether low-SES parents, due to having less accurate beliefs, make larger investment mistakes. To see how to test this hypothesis, assume there are two types of parents, low-accuracy and high-accuracy, who are identical except that the beliefs of low-accuracy parents have a lower correlation with true performance than the beliefs of high-accuracy parents.

By equation (4), if $\text{std}(\alpha)/\text{std}(a)$ does not vary by parent type (which it does not in the data used in this paper), then the slope of beliefs on true performance would

¹¹ Beliefs about academic potential beyond a should be more important if parents think a is a noisy measure of true potential. In my empirical setting, as discussed in Section IIIA, noise in the a measure does not appear to play a large role, perhaps because the a measure averages across multiple tests.

be flatter for low-accuracy parents than for high-accuracy parents.¹² As a result, by Prediction 1, the following empirical pattern would suggest that low-accuracy parents' inaccurate beliefs have a larger impact on their investments: the slope of the actual investment function would be flatter among low-accuracy than high-accuracy parents at baseline. To test for this heterogeneous baseline attenuation, I use Prediction 2 and test whether providing information increases the slope of investments on performance more for low-accuracy parents than for high-accuracy parents. Note that this test is only appropriate for investments whose s^* functions are relatively homogeneous by parent type.

Uncertainty.—The earlier statement that a parent with mean beliefs α_{ij} would choose inputs $s^*(\alpha_{ij})$ depends on the assumption that the utility-maximizing choice does not depend on her beliefs uncertainty, σ_{ij} . While this would be the case in some settings (e.g., with quadratic loss), in others, having higher uncertainty could cause utility-maximizing investments to diverge from $s^*(\alpha_{ij})$. Denote chosen investments in the general case as

$$(5) \quad s^{**}(\alpha_{ij}, \sigma_{ij}) = \operatorname{argmax}_s \int U(s, \alpha_{ij}) g(\alpha_{ij}, \sigma_{ij}) d\alpha_{ij}.$$

The slope of s^{**} with respect to mean beliefs, $\partial s^{**} / \partial \alpha$, could depend on σ_{ij} . Although this is not definitive, one might expect that higher belief uncertainty would decrease the magnitude of the slope ($|\partial s^{**} / \partial \alpha|$), thus further attenuating the slope of investments on a relative to the case where there is no uncertainty. The intuition is that uncertainty may make parents hesitate to rely as strongly on their mean beliefs when making their investments. Appendix Section B presents one potential framework yielding this prediction. I test for this effect empirically.

II. Context and Experimental Design

The setting for the experiment is Malawi. Overall, the education system in Malawi is similar to the systems in much of sub-Saharan Africa and other developing countries, in terms of the information given to parents and the overall structure. Primary school in Malawi covers grades 1–8. Although it has technically been free since 1994, it does involve expenditures. Parents in the study sample spent an average of 1,750 Malawi Kwacha (MWK) annually per child, roughly US\$10.6 or 1.6 percent of annual household income. The main expenditures are uniforms (33 percent), informal but required school fees (22 percent), and supplemental investments such as school supplies, tutoring, and books (45 percent). The access rate to the first grade of primary school is above 95 percent, but dropouts are common. Sources vary, but all suggest the primary school completion rate (conditional on enrolling) is less than

¹² An alternate way to express that one group has greater “belief inaccuracy” would be to say that the group has higher $E[(a - \alpha)^2]$. To map that to attenuation, express the OLS formula for the slope of α on a as $\frac{1}{2} + \frac{1}{2\operatorname{var}(a)}(-E[(a - \alpha)^2] + [Ea - E\alpha]^2 + \operatorname{var}(\alpha))$. This formulation makes clear that beliefs will be more attenuated for the group with higher $E[(a - \alpha)^2]$ if the following conditions are met, as they are in the data used in this paper: (i) the variances of beliefs and of true performance are similar across groups, and (ii) the bias in population-level beliefs ($Ea - E\alpha$) is similar across groups (see online Appendix Table C.1).

60 percent (World Bank 2010). Secondary school, covering grades 9–12, is not free; annual fees for government secondary schools range from 5,000–10,000 MWK per year (US\$30–\$60, over four times the median primary-school expenditures in the sample) (World Bank 2010). Uniforms and supplementary supplies are additional expenses. Many children do not attend because of the high costs. Secondary slots are also limited, with admissions governed by an achievement test administered at the end of primary school.

As in many other countries, schools are required to send report cards home each term with average achievement test scores; all schools in the sample for this study comply with the rule. The reports vary by school, but all are required to include average absolute test scores and the corresponding grade on the standard Malawian grading scale of 1–4. (Online Appendix D contains an example from the study sample.) However, the official report cards are often hard for parents to understand, or do not reach them at all. According to baseline survey data summarized in online Appendix Table C.2, 60 percent of parents state that they do not know their child's performance from the last report. Among that 60 percent, 50 percent did not receive the report card at all. Since students are supposed to deliver the reports, children could either lose or choose not to deliver them: parents of students who performed poorly are less likely to receive the report. Among the remaining parents who did not know their children's performance, a key reason seems to be an inability to read or understand the report: 50 percent of those parents are illiterate and 70 percent do not know basic details about the report card's structure (i.e., do not know at least one of whether the report card contained grades, positions, or scores (column 4)).

Report card knowledge is heterogeneous by parents' education. The rate of not knowing their children's performance is 20 percentage points (pp) higher among parents with below-median education than above-median, even conditional on school fixed effects. The reasons for lack of knowledge also differ by parental education. Not receiving the report card is a more common explanation among more-educated parents (even though, in absolute terms, more high-educated parents receive the report cards). In contrast, for less-educated parents, failure to understand the report card is more prevalent, with a higher rate of not understanding the basic report card structure.

A. Experimental Design

The experiment delivers academic performance information to randomly selected parents and measures the effects on educational investments and decisions. Although the school report cards should ostensibly already deliver this information, the report card system does not always succeed in conveying the information; the experiment presents the information more clearly. To fit the framework presented in Section I, the experiment should provide information about the individual-level trait on which parents' educational investments depend. In qualitative interviews, we asked parents what information is most helpful for making decisions about their children's education; academic performance (i.e., scores on school-administered exams) was the nearly universal response. If parents were wrong about the education production function, a second objective relevant for welfare would be to use the trait most correlated with *actual* individual-level returns. Academic performance also likely

meets this second objective: it determines progression through school and selection into secondary school, thereby almost surely affecting the returns to investment. “Innate” ability is another possible determinant of returns, but, as has been extensively documented, it is difficult to measure “innate” ability; any measure would represent some combination of innate ability and past inputs. (See Section III E for further discussion.)

Sample Selection.—The study worked with 39 schools in two districts (Machinga and Balaka) in Malawi. Schools were selected randomly from the universe of primary schools, oversampling schools with high and low expected levels of parent education to increase heterogeneity in parent education within the sample. The study team first conducted a census at schools, mapping the sibling structures for all students in grades 2–6; these grades were chosen because they span most of primary school. Since one of the outcomes to be examined is inter-sibling trade-offs, multiple-sibling households were used as the sampling frame (fewer than 3 percent of the households in Malawi who have children and have completed their fertility have only one child). The team also gathered achievement test data from the most recent term (term 2 of the 2011–2012 school year) for use in the intervention.

Based on the test score and sibling data, a sample of 3,451 households with at least two children enrolled in grades 2–6 with test score data was drawn. For households with more than two children, two children were randomly selected. Because one inclusion criterion was that children needed test score data, students who have the highest absence rates (and whose parents might have the largest information problems) are underrepresented in my sample.

Randomization.—I randomly assigned one-half of the households in the sample to a treatment group that received information about their children’s test scores, and one-half to a control group which did not.¹³ The randomization was stratified on a test score measure (between-sibling score gap), and a proxy for parent education (the estimated literacy rate in the household’s village), since one *ex ante* goal was to look at heterogeneity by parent education.

Eligibility Interviews.—Sample selection and randomization were based on data gathered from students at school and on school administrative data. Household eligibility (i.e., whether both siblings lived in the household and were still enrolled in school) was then verified through an eligibility questionnaire with parents. Among the 3,451 sampled households, 21 percent 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 percent (2,634 households) were located at their homes, available, and consented to participate in the baseline survey. Thus, the final experimental sample comprises 5,268 parent-child pairs. Both eligibility and baseline survey completion are unrelated to treatment assignment.

¹³ Half of the treatment group was also assigned to receive an add-on intervention designed to test a hypothesis intended for study in separate work: that providing more detailed information would increase parental engagement. This group received additional skills information (e.g., whether their child could add three-digit numbers; see online Appendix E for sample). In this paper, I ignore this add-on treatment and pool the treatment households. I do not find that this treatment had an effect on the pre-specified outcomes.

Baseline Survey Visit.—Surveyors visited all sampled households and asked to speak with the parent who is the primary decision-maker about education.¹⁴ Surveyors then conducted a baseline survey, which included a module on education spending and beliefs about children's test scores. While eliciting baseline beliefs about test scores, surveyors explained the grading scale used by schools to parents; they also reviewed a sample report card which had the same format as those later delivered to the treatment group. This was done to aid the elicitation of beliefs and to hold knowledge of the grading scale and report card format constant across treatment groups. After the survey, during the same visit, surveyors conducted the information intervention for the treatment group.

Information Intervention and Report Cards (Treatment Group Only).—Surveyors walked parents through two report cards (one per child) describing the academic performance of their children. The order was randomized. The reports showed children's performance on all tests administered in the most recent school term, specifically: the percent score (an absolute measure), the corresponding grade on the Malawian grading scale, and the within-class percentile ranking (see online Appendix F for more details). The statistics were listed for the three subjects that Malawian educators deem most important (math, English, and Chichewa, the local language) and for "overall" (the average of the three). The report card also showed the number of individual tests included in the averages; teachers conducted tests 4.5 times per term on average, with each individual subject tested 4.2 times (not all subjects were tested every time). The correlation between scores on different tests within the term is roughly 0.8 for overall performance, and 0.6–0.8 within subjects.

A sample report card is presented in online Appendix G. The format was chosen based on a series of focus groups; the primary selection criterion was whether uneducated parents could understand it. Surveyors, who were trained on how to explain the information clearly, walked treatment parents through every number on the report cards.

B. Data and Outcomes

The analysis uses several data sources, including data from surveys with parents and administrative data from schools. Online Appendix F.1 shows the survey questions that measured beliefs and the experimental outcomes, and the scripts for the information intervention.

(1) *Baseline Survey Data:* The baseline survey was rolled out immediately after term 2 of the school year, which ended in March 2012, and ran from April to June 2012. The survey included modules on demographics, education spending, and the perceived returns to education. Mean beliefs about academic performance were measured by asking parents about the same performance metrics that were later delivered in the intervention report cards: average scores and percentile rankings on the term 2 school exams in math, English, Chichewa, and overall. We used the

¹⁴If that parent was unavailable, the surveyor spoke with the second parent if he/she was present and knowledgeable about educational decisions. If not, the surveyor returned later.

same measure later used in the intervention so that any gaps between believed and true performance represent belief inaccuracies, not differences between measures. Beliefs uncertainty was measured by asking parents to distribute tokens across bins representing score ranges (e.g., 0–20).

(2) *First Endline Survey: Endline Beliefs and Experimental Outcomes:* This survey was conducted immediately after the baseline survey and information intervention; see Figure 2 for the data collection time line. This was done for budgetary reasons, but does have the advantage that the outcomes were measured before parents had a chance to speak with others, allowing the outcomes to more cleanly reflect parents' preferences, as opposed to the preferences of the people they talk to, including their children. Recall that there are two main categories of investment outcomes: (i) a series of real-stakes investment options and decisions presented to parents through the experiment ("experimental outcomes"); and (ii) more traditional endline outcomes measured during the year after the intervention, such as enrollment ("non-experimental" or longer-term outcomes). This survey measured the first category, which is described in more detail in Section IIB, as well as endline beliefs. When measuring endline beliefs, because I wish to assess whether information affects the beliefs underlying parents' behavior, I want to know both whether (a) parents understood and believed the information presented in the intervention, and (b) the information is relevant for their decisions going forward. As a result, surveyors asked parents what score they thought their child would receive if he took an exam *that same day*. Asking about the previous-term scores as done in the baseline survey would only have measured (a), since those exams happened in the past; asking about a (hypothetical) same-day measure allows us to also assess (b). I refer to these beliefs as "endline beliefs." Beliefs uncertainty was not measured at endline due to budget constraints (the process of measuring beliefs uncertainty was lengthy).

(3) *Non-Experimental (Longer-Term) Outcomes:* I also collected two types of data in the year following the intervention: (i) information from a second endline survey of parents one year after the intervention (June–July 2013), which I use to examine treatment effects on dropouts and expenditures; and (ii) administrative data on attendance gathered roughly one month after the intervention (July 2012). These outcomes allow me to establish the policy relevance of the findings. For the one-year second endline data collection, given the very limited budget, I focused on outcomes where (a) I expected results and (b) data collection costs were lower. I thus focused on dropouts and expenditures, rather than academic performance. Dropouts and expenditures are parental decisions that are easy to adjust, whereas academic performance reflects many other factors. There was sufficient budget to include roughly 900 households in the second endline survey sample. Of the households selected for the sample, 98 percent (893) were successfully surveyed, balanced across treatment group. The administrative attendance data were gathered by giving schools a template to record the data for the month following the intervention, and were collected from 35 percent of the sample. During the attendance data collection, we were able to collect data on endline exams for 9 percent of the sample; this allows me to validate the accuracy of the baseline academic performance measure, but does not give




FIGURE 2. OVERVIEW OF DATA COLLECTION

Notes: For any given household, all Day 1 activities were conducted on the same day as the baseline survey; across the sample, the baseline survey was rolled out over the course of two months.

sufficient statistical power to look at treatment effects. See online Appendix F.2 for more detail on the sample and data for the non-experimental outcomes.

Experimental Outcomes.—During the first endline survey, surveyors presented parents with three real-stakes investment decisions (the “experimental outcomes”). While the non-experimental outcomes are preferable from a policy-relevance perspective, these experimental outcomes have several advantages. Primarily, they enable me to include outcomes that are designed to have clear predictions for how the efficient investment depends on student performance, allowing for a clean test of whether misinformation causes mistakes. They also allow me to include outcomes where the “preferred investment function” is homogeneous by parental education, allowing me to cleanly test for whom inaccurate beliefs matter more. Finally, they are very precise, allowing for detailed heterogeneity analysis.

The outcomes include a combination of monetary investments involving cash outlays, and choices between free options. The latter allows us to abstract away from credit constraints, which is useful for heterogeneity analysis as it enables cross-household comparisons that are not confounded by household wealth. The outcomes also incorporate both smaller investments that primarily assess effects on the *types* of investments parents choose (i.e., whether they tailor the input mix correctly to their children’s needs) and a larger outcome designed to detect treatment effects on the *level* of investment across children.

Outcomes Capturing the Type of Investment: The first decision presented to parents is a choice among free grade-specific workbooks with different difficulty levels.

We gave parents four free books: an English and a math book for each of their two children. For each book, parents were allowed to choose between three levels of difficulty: beginner, average, or advanced. The obvious prediction is that book difficulty choice will increase in perceived performance. The second investment is the willingness to pay (WTP) for grade-specific, subject-specific remedial textbooks in math and English. WTP was evaluated using a Becker-DeGroot-Marschak (BDM) methodology, which gives respondents an incentive to report truthfully (see online Appendix H for description). The elicitation was real-stakes, with parents paying out-of-pocket for the textbooks using their own money, and the maximum price on the price list equal to the full market price. The average WTP for a book was substantial: 324 MWK (\$2), or roughly 20 percent of mean annual per-child educational expenditures. We use remedial textbooks (textbooks perceived by teachers as substitutes with performance). Thus, the prediction is that WTP will be higher for the subject in which parents think their child is doing worse; this was an *ex ante* prediction, later confirmed by baseline survey questions showing that 95 percent of parents believe the textbooks are substitutes with performance.

Both the textbooks and free workbooks have clear predictions for parents' beliefs about the "right choice" (i.e., the perceived production function). An additional advantage is that both have clear predictions for the *actual* right choice and true production function. For example, the advanced workbook was designed specifically to be better for the higher performers in the sample. This enables a stronger argument that parental mistakes due to misinformation about child performance lower actual (not just perceived) returns.

Outcome Capturing the *Level* of Investment across Children: Secondary schooling is the first high-cost educational investment in Malawi. Few parents in the sample could afford school fees for all of their children; many cannot pay for a single child. My third investment introduces a short-run, real-stakes proxy for secondary schooling. We conduct a lottery, in which the prize is four years of government secondary school fees for one child in every 100 households (worth roughly US\$120–\$240 at the time of the experiment). Parents were given nine tickets for the lottery and were asked to allocate the tickets across their two children. There are many "binary" choices in education where credit-constrained parents must choose between a lumpy investment in one child or the other; for example, if parents can only afford to send one child to secondary school or college. The lottery ticket allocation—and in particular, which child the parent allocates more tickets to—was designed to proxy for these types of decisions.¹⁵

There are two main channels through which academic performance would affect the expected return of a lottery ticket. First, through the earnings return to secondary school: 95 percent (78 percent) of parents believe that secondary school increases the earnings of higher-performing students weakly (strictly) more than the earnings of lower-performing students, and, on average, parents perceive the earnings

¹⁵ Although a single ticket could have also accomplished this goal, I used multiple tickets to increase the power to detect small shifts and to allow me to make use of this lottery in a separate paper studying inequality aversion. As expected, most parents (74 percent) split their nine tickets as evenly as possible, consistent with an aversion to inequality between their children. Thus, in most cases, the analysis reduces to which child the parents give their ninth ticket to, which proxies for the child they would choose in a binary choice.

TABLE 1—BASELINE SUMMARY STATISTICS

	Full sample		Control Mean	Treat Mean	Treat – Control		
	Mean	SD			Mean	SE	p-val T=C
<i>Panel A. Respondent background</i>							
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
<i>Panel B. Household background</i>							
Family size (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
<i>Panel C. Student information</i>							
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. ^b	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
<i>Panel D. Academic performance (average achievement scores)</i>							
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	51.2	22.5	51.5	51.0	−0.55	0.59	0.35
(Math – English) score	0.71	19.5	0.93	0.5	−0.53	0.51	0.3
<i>Panel E. Respondent's beliefs about child's academic performance</i>							
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
SD of individual beliefs about score	7.69	10.1	8.08	7.28	−0.8	0.38	0.03
<i>Panel F. Gaps between believed and true academic performance</i>							
abs[believed – true overall score]	20.4	14.5	20.4	20.3	−0.12	0.43	0.78
abs[believed – true math score]	25.8	18.0	25.8	25.7	−0.1	0.52	0.85
abs[believed – true English score]	21.4	16.4	21.6	21.1	−0.57	0.48	0.23
abs[believed – true Chichewa score]	23.8	17.5	23.7	23.9	0.19	0.51	0.72
abs[believed – true (math – English) score]	22.1	17.4	22.3	21.9	−0.44	0.51	0.39
abs[believed – true overall score (child 1–2)]	18.7	15.1	18.9	18.5	−0.34	0.59	0.56
Believed – true overall score	15.6	19.5	15.6	15.6	−0.08	0.58	0.89
Believed score higher than true score	0.79	0.41	0.79	0.79	0.01	0.01	0.65
Wrong about who (child 1 or 2) is higher-scoring	0.31	0.46	0.32	0.31	−0.01	0.02	0.63
<i>Panel G. Beliefs about complementarity</i>							
Believes educ. and achievement complementary ^c	0.91	0.29	0.9	0.91	0	0.01	0.68
<i>Sample sizes</i>							
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 test of equality across treatment and control 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 higher-performing child would increase “more” or “much more” than the earnings of a lower-performing child from getting a secondary education.

increase to be 90 percent higher for a hypothetical child in the top decile of performance than for one in the bottom decile. Second, since admissions is governed by performance on a standardized achievement test, the probability of admission to secondary school increases with performance, a fact that 98 percent of parents are aware of. Thus, the (perceived) expected value of the fees paid and the probability of attending both increase with performance. Taking both channels together, a back-of-the-envelope calculation based on parents' beliefs suggests that the perceived return is over 300 percent higher for students in the top versus bottom performance decile.¹⁶ Thus, the prediction is that parents will allocate more lottery tickets to higher-performing children.

C. Summary Statistics and Balance

Table 1 presents summary statistics and tests for balance across the treatment and control groups. Seventy-seven percent of respondents are female, and 92 percent are the primary education decision maker in the household. Average levels of parental education are low, at 4.7 years. Households are large, with an average of 5 children. Sampled children were 12 years old on average, primarily aged 8 to 16, and 51 percent female. To test balance, I regress each variable on a dummy for being in the treatment group. The differences between the treatment and control groups are never large, with a joint test of equality failing to reject the null that all are 0 (p -value 0.67). Only 1 of the 39 variables is statistically significant at the 5 percent level: baseline math scores. To ensure this imbalance does not affect the results, all regressions control for an academic performance measure, although the results are robust to omitting this control.

III. Empirical Results



I begin by showing that parents have inaccurate beliefs about their children's academic performance. I then demonstrate, using first the experimental and then the non-experimental outcomes, that their belief inaccuracies cause them to make mistakes when making decisions. Finally, I provide evidence linking information frictions with poverty, and discuss the implications of these frictions for welfare and the average level of investment in education.

A. Beliefs


RESULT A1: *Parents' beliefs about academic performance are inaccurate.*

Data from the baseline survey can be used to assess the accuracy of parents' beliefs about their children's "academic performance," i.e., scores on school-administered exams the prior term. Panel A of Figure 3 presents the average of the absolute value of the gap between parents' mean beliefs about their children's academic performance and their children's true academic performance. Scores are absolute

¹⁶See online Appendix I for calculation.

Panel A. Gap between true test scores last term and *baseline* beliefs about scores last termPanel C. Gap between true test scores last term and *endline* beliefs about likely score on a hypothetical test

Panel B. Attenuation of baseline beliefs



Panel D. Attenuation of endline beliefs




FIGURE 3. BELIEFS RESULTS

Notes: Data sources are survey data and administrative baseline test score data. Scores are absolute percentages, expressed on a scale from 0 to 100. All scores and beliefs are about overall (as opposed to subject-specific) performance. Panel A displays the average absolute value of the gap between children's true test scores last term and parents' beliefs (measured at *baseline*) about these test scores; it shows that inaccuracies are large, and balanced across the control and treatment groups. Panel B shows attenuation in baseline beliefs by plotting locally linear regression lines with beliefs about last-term test scores as the dependent variable and true test scores as the *x*-axis; it shows that baseline beliefs are attenuated (i.e., that the slope is less than 1 and so they do not move 1-to-1 with true scores), and that this is balanced across the treatment and control groups. Panel C displays the average absolute value of the gap between children's last-term true test scores and parents' beliefs (measured at *endline*) about their children's performance on a hypothetical test taken that same day; it shows that information moves parents' beliefs toward the signal. Panel D shows attenuation in endline beliefs by plotting locally linear regression lines with beliefs (measured during the first endline survey) about performance on a hypothetical test as the dependent variable and last-term true test scores as the *x*-axis; it shows that information decreases the attenuation.

percentages, expressed on a scale from 0 to 100.¹⁷ The graph shows the treatment

¹⁷ The online Appendix shows that the main results are robust to using relative performance (within-class percentiles) instead of absolute performance (Tables C.3, C.4, and C.5). In my sample, absolute and relative performance are highly correlated (0.8). Rogers and Feller (2016) compare the effects of relative versus absolute information about student absence rates from school; such a comparison was not a goal of this study. That being said, online Appendix C.24 shows results when both measures are analyzed simultaneously; parents seem to respond more to

and control groups separately to demonstrate baseline balance.

The average gap is large: 20 points, or 1.2 standard deviations of the performance distribution for overall performance. Panel F of Table 1 shows that mean beliefs about individual subjects like math, between-subject performance (math versus English), and the between-sibling gap are also inaccurate. Beliefs about the between-sibling gap diverge from the true gap by 1.1 standard deviations on average, with 31 percent of parents wrong about which child is higher-scoring. While parents overestimate on average, 21 percent of parents do not.

As described in Section I, these belief inaccuracies should cause mean beliefs to not move 1-to-1 with true scores and instead to have a slope less than 1 on true scores. Panel B of Figure 3 substantiates this for overall performance with a local linear regression of mean beliefs on true performance: the slope is visually less than 1. This attenuation in the slope captures the fact that the correlations between believed and true performance are low: 0.3 for overall performance, as depicted in the graph, and 0.2–0.3 for performance in the individual subjects like math. Since these tests determine progression through school, these inaccuracies are likely relevant for a broad range of investments.

One natural question is whether these “inaccuracies” in beliefs simply reflect noise in the performance measure. The data suggest otherwise. The correlation between tests taken during the term is 0.8 for overall performance, and 0.6–0.7 within subjects, which suggest high test reliability; these correlations are notably higher than the correlations between parents’ beliefs and the term-average scores (0.2–0.3). I also have data on *future* test scores for a small subset of the sample which shows that baseline test scores are nine times more predictive of future test scores than parents’ baseline beliefs are.¹⁸ Moreover, we can use the experiment itself to provide more evidence on this issue: If providing information to parents causes them to update their beliefs, it suggests that parents themselves believe that there is additional meaningful content in the information that was not reflected in their baseline beliefs.

RESULT A2: *Providing information aligns beliefs better with students’ test scores.*

I now examine whether information changes beliefs and decreases attenuation by looking at the impact of information on mean beliefs measured at endline. Recall that, unlike beliefs measured at baseline, the beliefs question asked at endline was not asking about last-term test scores; instead, it asked how well parents thought their child would do on a hypothetical test taken that same day. The prediction is thus that providing information should decrease the gap between parents’ endline beliefs and their child’s last-term scores, as their posterior beliefs move in the direction of the signal; the gap, however, should not fall to 0, unless parents place no weight on other factors (e.g., their assessment of their child’s recent progress). Panel C of Figure 3 graphs the absolute value of the gap between true baseline (last-term) performance

absolute than to relative information. Online Appendix F.4 discusses this in more detail, as well as explaining the other reasons that the analysis uses absolute performance.

¹⁸See online Appendix Table C.6 for a regression using control group data: the coefficients on current test scores and beliefs are 0.74 and 0.08, respectively. Misunderstanding the difficulty of the grading scale also does not drive the results: the patterns are similar for within-class percentile ranks (online Appendix Table C.3).

and endline beliefs, separately by treatment group. Information cuts the gap nearly in half. Panel D of Figure 3 shows that attenuation correspondingly decreases: the slope of endline beliefs on true baseline scores is over twice as steep for the treatment group as for the control group.

I formally test whether information increases the slope of endline beliefs by running the following regression:

$$(6) \quad y_{ij} = c_0 + c_1 A_{ij} \times Treat_i + c_2 A_{ij} + c_3 Treat_i + c_4' X_{ij} + \varepsilon_{ij},$$

where i indexes households, j indexes siblings, y_{ij} is the parent's endline beliefs about overall academic performance, A_{ij} is baseline overall academic performance, $Treat_i$ is an indicator for being assigned to the treatment group, and X_{ij} is a vector of control variables.¹⁹ Standard errors are clustered at the household level. Column 1 of Table 2 confirms the increase in slope ($c_1 > 0$). Information thus has a statistically significant "first-stage" effect on beliefs, allowing us to use the experiment to examine the effects of information and beliefs on investments.

B. Results: Experimental Outcomes

I first examine the "experimental outcomes," i.e., the investment decisions presented to parents in the first endline survey, which allow me to cleanly document investment mistakes, before turning to the "non-experimental outcomes" to establish greater policy relevance. I begin by using data from the control group to provide motivating evidence of mistakes, and then present the information treatment effects.

RESULT B1: *Control group parents attempt to tailor their investments to performance, but partly fail.*

Data from the control group can be used to study how baseline parental investments depend on parents' beliefs about performance. This can give us insight into parents' preferred investment function and the likely production function that they have in mind. We can then compare this with how investments vary with *true* performance. A divergence between the two relationships would suggest that inaccurate beliefs may affect investments.

Figure 4 compares the preferred investment function (investments plotted against believed performance: the dashed lines) with the actual investment function (investments plotted against true performance: the solid lines). Note that the y-axes for both lines represent investments, but the x-axes differ. Both are locally linear regressions using control group data. I first interpret the preferred functions and then compare them to the actual.

Panels A1 and A2 of Figure 4 presents the results for math and English workbook difficulty choices graphically. Recall that for each book, parents could choose from

¹⁹Results are robust to excluding the controls (see online Appendix Tables C.7 and C.8). Controls include school fixed effects (FE), the between-child score gap, average parental years of education, a parental education proxy used for stratification, child and parent gender, and grade FE. This includes all variables underlying the stratification but not the stratum FE themselves as some strata are very small.

TABLE 2—EXPERIMENTAL OUTCOMES: INFORMATION TREATMENT EFFECTS ON THE SLOPE OF INVESTMENTS ON ACADEMIC PERFORMANCE

	Endline beliefs (1)	Math workbook difficulty level (2)	English workbook difficulty level (3)	ln(math textbook WTP) – ln(English textbook WTP) (4)	Secondary school lottery tickets (5)
Treat × Score	0.40 [0.025]	1.33 [0.093]	1.25 [0.096]	0.013 [0.0022]	0.036 [0.0052]
Score	0.32 [0.018]	0.65 [0.066]	0.76 [0.073]	0.0023 [0.0016]	0.015 [0.0051]
Treat	−25.9 [1.33]	−91.0 [4.91]	−68.4 [4.83]	0.14 [0.041]	
Observations	5,244	5,239	5,239	5,219	5,258
R ²	0.372	0.218	0.206	0.036	0.154
Score used	Overall	Math	English	English – math	Overall
Household FE	No	No	No	No	Yes

Notes: Data sources are baseline survey, baseline test score data, both endline surveys, and endline administrative data. Each observation is a child. Standard errors are clustered at the household level. Workbook difficulty choices are coded as 0 for beginner, 100 for average, and 200 for advanced. The dependent variable in Column 1 corresponds to the parent's endline beliefs about the child's overall score on a hypothetical test taken the same day as the endline survey.

The regressions test whether information changes the slope of investments on children's academic performance (where academic performance is measured as children's average scores on school-administered achievement exams). One way to interpret the results is to compare the baseline slope in the control group (coefficient on *Score*) with the increase in the slope in the treatment group (coefficient on *Treat* × *Score*) to see how much the slope has increased as a result of information. Take for example column 2. The ratio of the coefficient on *Treat* × *Score* (1.33) to the coefficient on *Score* (0.65) shows us that the slope has increased by roughly 200% (1.33/0.65), so that the treatment slope is roughly 3 times as large as the control slope. The rough interpretation of the slope in the control group for that column is that, if the child's math score increases by 1 point, the chance that her parent chooses the next higher difficulty level of the free book increases by 0.65%. Regressions control for school FE, average parental years of education, a parental education proxy used for stratification, the between-child score gap, child gender, grade FE, and parent gender; column 5 also has a household FE.


three different difficulty levels. The y-axis represents the chosen difficulty level, with the three different levels parametrized as 0/1/2 for simplicity, but the results are robust to other parametrizations.²⁰ I focus first on the dashed lines, which represent parents' preferred choice given their beliefs about their child's math or English score, represented by the x-axis. The obvious prediction is that book difficulty choice should increase in believed performance, and consistent with this prediction, the dashed lines for both English and math slope steeply upward.

Panel B of Figure 4 presents similar results for the second investment, the willingness to pay (WTP) for subject-specific textbooks in math and English. Because the textbooks are remedial, the prediction is that WTP will be higher for the subject in which parents think their child is doing worse. The use of the between-subject WTP (math–English) holds constant other factors, such as the child's overall performance, which is advantageous for this test as it provides clean predictions.²¹


²⁰ See online Appendix Tables C.9 and C.10 for robustness to an ordered probit specification.

²¹ The ex ante plan was to analyze within-child, between-subject WTP, since surveys with teachers showed that teachers universally think that, within child, the remedial textbooks have higher returns in a *subject* in which the child is behind. In contrast, teachers have more mixed opinions about whether remedial textbooks have higher returns for a *child* who is behind in a subject relative to a child who is ahead. These mixed opinions reflect the fact that a child who is behind in one subject also has lower performance on average, and so might be less inclined to


Panel A1. Difficulty level chosen for free math workbooks



Panel A2. Difficulty level chosen for free English workbooks



Panel B. WTP for remedial textbooks



Panel C. Secondary school lottery




FIGURE 4. IN THE CONTROL GROUP, THE SLOPE OF INVESTMENTS ON TRUE ACADEMIC PERFORMANCE IS ATTENUATED RELATIVE TO THE SLOPE ON BELIEVED PERFORMANCE

Notes: Control group data only. Data sources are survey data and administrative baseline test score data. Lines are locally linear regression lines with investments as the dependent variable and either true (solid line) or believed (dashed line) baseline academic performance as the x-axis. For the workbook graphs (panels A1 and A2), the dependent variable is the parent's choice of difficulty for a free workbook, where 0 corresponds to the beginner workbook, 1 corresponds to the average, and 2 to the 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. Because the textbooks are remedial, the prediction is that this should increase in the child's English relative to math performance. For the secondary school lottery (panel C), the dependent variable is the number of secondary school lottery tickets given to the older relative to younger child in the household, and the believed score gap is the gap in parents' beliefs about their children's overall test scores. The gray areas are 95 percent confidence intervals.

use a textbook, even a remedial one, since he/she may be more likely to drop out of school, etc. The within-child, between-subject comparison holds those factors constant. See online Appendix Table C.11 for the results for math and English textbooks separately.

In panel B, as in panels A1 and A2, the dashed lines are the preferred investment lines. The x -axis shows beliefs about performance in English relative to math. The y -axis shows the log of WTP for the math textbook minus the log of WTP for the English textbook.²² For presentation purposes, English is flipped relative to math on the y -axis; the prediction thus becomes that the line will have a positive slope. The dashed line slopes steeply upward, consistent with the prediction that WTP increases the further behind a child is in a given subject.

Panel C of Figure 4 shows the secondary school lottery ticket allocation results. The dashed line plots the difference in tickets allocated to the older versus the younger child in the pair, with the x -axis the gap in perceived scores between the older and younger child.²³ Consistent with the *ex ante* prediction that allocating more tickets to higher-performers yields higher returns, the line slopes upward: parents give more tickets to the child they think is higher-performing.

I now compare the slope of the preferred investment functions just discussed with the slopes of the *actual* investment functions, depicted by the solid lines in Figure 4. The solid lines have the same y -axes as the dashed lines, but different x -axes: their x -axes are true performance instead of believed. The prediction is that, if parents base their investments on their inaccurate beliefs, then the slope of their investments on true performance will be attenuated relative to the slope on beliefs. And, in fact, the graphs show precisely this pattern: the slopes on true performance are only 15–40 percent as large as the slopes on beliefs. This suggests that parents try to tailor their investments to performance, but that their inaccurate beliefs prevent them from doing so. Since returns depend on *true* performance, if parents *knew* that, say, their child had a math score of 80, they would choose the highest difficulty book for him, but many parents do not know that and so fail to choose their preferred option. This evidence is suggestive, however, not causal; both beliefs and performance could be correlated with other factors affecting investments. An experiment, in contrast, can establish causality: I can test whether information undoes the attenuation. I turn to this next.

RESULT B2: *Information substantially increases the slope of investments.*

I now use the information experiment to test whether information increases the slope of investments on actual performance. Figure 5 shows locally linear regressions of investments on true performance for the treatment group (dashed line) and control group (solid line). Both lines have true performance as the x -axis. Note that Figure 4 examines only the control group, and thus the solid lines in Figures 4 and 5 are identical, as they depict the same data.

The figures confirm that information frictions cause mistakes: for all three investments considered, the information treatment substantially increases the slope of the investment functions. I perform a formal test of the change in slope by estimating

²² Only 6 percent of observations are 0s, which I replace with the log of 10 percent of the lowest price list value. The results are robust to using other values (e.g., 50 percent), dropping the 0s, or using levels instead of logs: see online Appendix Table C.12.

²³ Since the lottery is a within-household allocation, to depict it graphically, we need to order the two children in some way. Parents identified age in focus groups as the second most important factor for investment (behind performance) so I order using age, but the graphs look similar with any order.




FIGURE 5. THE INFORMATION TREATMENT INCREASES THE SLOPE OF INVESTMENTS ON TRUE ACADEMIC PERFORMANCE

Notes: Data sources are survey data and administrative baseline test score data. Lines are locally linear regression lines with investments as the dependent variable and either true (solid line) or believed (dashed line) baseline academic performance as the x-axis. For the workbook graphs (panels A1 and A2), the dependent variable is the parent's choice of difficulty for a free workbook, where 0 corresponds to the beginner workbook, 1 corresponds to the average, and 2 to the 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. Because the textbooks are remedial, the prediction is that this should increase in the child's English relative to math performance. For the secondary school lottery (panel C), the dependent variable is the number of secondary school lottery tickets given to the older relative to younger child in the household, and the believed score gap is the gap in parents' beliefs about their children's overall test scores. The gray areas are 95 percent confidence intervals.

equation (6) using the experimental outcomes as the outcome variables and using the relevant academic performance metric as A_{ij} (e.g., math for math workbooks). The prediction is that the information treatment makes the slope steeper, so that $c_1 > 0$ (with c_1 the coefficient on $A_{ij} \times \text{Treat}_i$). The key prediction regards c_1 ; c_3 , the coefficient on Treat_i , is not particularly meaningful as it is just driven by the scaling of

the A_{ij} variable, representing the treatment effect for those for whom $A_{ij} = 0$ for the particular A_{ij} measure used in that regression. For example, for the textbook regression, it is the treatment effect for those who have the same performance on math and English (i.e., $math - English = 0$).

Columns 2 through 5 of Table 2 present the results for the math and English workbook difficulty choices; the log of WTP for the math textbook minus the log of WTP for the English textbook; and the secondary school lottery tickets received. Since secondary school lottery tickets are inherently a within-household allocation (one child's allocation fully determines the other's), the lottery regression is estimated with a household fixed effect. Consistent with the graphical evidence, across all outcomes, c_1 is positive and statistically significant. The magnitudes are 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), we see that information causes investments to become 3–6 *times* more steeply aligned with performance across the various investments, i.e., the slopes increase by roughly 200–500 percent. This suggests that parents were making substantial mistakes at baseline.

One question is whether the treatment effects are driven by information increasing the salience of education. If salience effects were uniform, they would affect the *level*, not slope, of investment. One could, however, be concerned that salience effects vary and are correlated with performance. Since salience would likely be a household-level effect or would be correlated with child observables, online Appendix Table C.13 allays this concern by showing robustness to the inclusion of household fixed effects and child-level controls interacted with treatment.²⁴ Treatment effects on longer-term outcomes can also mitigate the concern, as salience may decrease over time. A second question is whether demand effects play a role here. The information treatment could cause the treatment group to align their investments more closely with the information delivered to them if they believe surveyors have that expectation. The use of real stakes for all investments, the standard approach to address demand effects, helps allay this concern. De Quidt, Haushofer, and Roth (2017) provide evidence that demand effects are modest with incentivized choices. Again, the treatment effects on longer-term outcomes can also help mitigate the concern, since these outcomes should not be subject to demand effects. I analyze these outcomes next.

C. Results: Non-Experimental (Longer-Term) Outcomes

The results above demonstrate that inaccurate beliefs cause parents to make mistakes when allocating educational investments. An open question, however, is the relevance for decisions outside of the experimental environment. I next turn to longer-run, non-experimental outcomes to show that information frictions are also relevant for outcomes that map more directly to the policy outcomes of interest.

²⁴ Relatedly, one could be concerned that providing information affects the salience of investing based on perceived performance. This concern is allayed by two facts: first, even investments in the control group are steeply sloped with perceived performance; and second, as I will show in Section IIID, the parents with the least accurate baseline beliefs experience the largest treatment effects.

TABLE 3—LONGER-TERM OUTCOMES: INFORMATION TREATMENT EFFECTS ON THE SLOPE OF INVESTMENTS ON ACADEMIC PERFORMANCE

	Enrollment (1)	ln(total educ. expenditures) (2)	Attendance rate (3)
<i>Panel A. Continuous versions</i>			
Treat × score	0.11 [0.038]	−0.0018 [0.0022]	0.021 [0.049]
Treat	−5.33 [2.10]	0.087 [0.11]	−1.36 [2.62]
Score	−0.018 [0.023]	0.0038 [0.0015]	−0.086 [0.034]
<i>Panel B. Binary versions</i>			
Treat × above-median score	3.77 [1.45]	−0.027 [0.074]	−0.12 [1.54]
Treat	−2.22 [1.15]	0.015 [0.061]	−0.17 [1.24]
Above-median score	−1.69 [1.38]	0.12 [0.069]	0.80 [1.38]
Observations	1,786	1,709	1,827
Control group mean	97.9	7.4	91.1
Score used	Overall	Overall	Overall

Notes: Data sources are baseline survey, baseline test score data, endline survey, and endline data collected from schools. Each observation is a child. Standard errors clustered at the household level. All regressions also control for grade FE, school FE, the between-child score gap, household-average years of parental education, child gender, parent gender, a parental education proxy used for stratification, and the baseline value of the dependent variable, if available (not available for enrollment). *Enrollment* is defined as being enrolled in school one year after the intervention. Both enrollment and attendance are scaled to be out of 100 (so enrollment, for example, is equal to 100 if the child is still enrolled and 0 otherwise). *Above-median score* means the child had an above-median baseline overall score.

However, the ex ante predictions for the preferred investment function are generally not as clear,²⁵ and precision is lower.

RESULT C1: *Information affects the slope of non-experimental investments.*

Here again, I examine the effect of information on the slope of investments. Panel A of Table 3 presents estimations of equation (6), all using overall scores as the performance measure. To aid in interpretation, panel B shows estimates using a binary score measure, specifically an indicator for whether a student has an above-median score. I consider three outcomes: primary school enrollment (dropouts), attendance, and expenditures. Of the three, primary school enrollment, which likely proxies most closely for the overall *level* of investment, provides the cleanest test: consistent with the literature, most parents believe additional years of schooling are more valuable for higher-performing children, whereas parental beliefs about the complementarity of expenditures or attendance with performance, as elicited in interviews, vary widely

²⁵It is also harder to use control group data to generate predictions for the production function parents have in mind; compared with the experimental outcomes, these outcomes have more omitted determinants, making the observational regressions harder to interpret. However, we can still use the information treatment effects themselves to infer the perceived complementarity/substitutability of the investments with performance.

across parents. The literature on attendance and expenditures is also limited, and there is little reason to expect the production function to be the same as for years of schooling. For example, conditional on having a child enrolled in school, parents may need to invest more in their lower-performing children to keep them on track.

Column 1 shows the primary school enrollment results. Consistent with the fact that nearly all parents believe years of schooling are a complement with academic performance, information increases the slope of the investment function. High-performing students in the treatment group are more likely to be enrolled in school one year later, while low-performing students are less likely to be enrolled. The change in the slope in panel A is significant at the 1 percent level.²⁶ Panel B shows that the magnitudes are economically meaningful. Among children whose parents found out they had above-median performance, dropout falls to nearly 0 percent (from a control group mean of 2 percent), whereas it roughly doubles for those with below-median performance, increasing from 2 percent to about 4 percent. These results highlight that information may not improve educational outcomes for all: it leads to reallocations, which can decrease investments for some. Since the literature suggests that schooling and ability are complements, these reallocations are consistent with an improvement in returns (Pitt, Rosenzweig, and Hassan 1990; Aizer and Cunha 2012). In contrast to the results for primary school enrollment, but perhaps to be expected given parents' heterogeneous beliefs regarding complementarity with performance, I find no significant effects for either expenditures or attendance. See online Appendix J for results on two additional secondary outcomes.

D. The Link between Information Frictions and Poverty

I now provide evidence linking information frictions with SES by testing for heterogeneity by a measure of SES: parental education. I use parental education because it is the least noisy measure of SES in the data and because limited education provides a potential channel for why low-SES parents are less able to access information.

RESULT D1: *Less-educated parents have less accurate beliefs.*

Panel A of Table 4 presents the results of the following regression which tests for heterogeneity in the attenuation of beliefs by parental education:

$$(7) \quad \tilde{A}_{ij} = d_0 + d_1 A_{ij} + d_2 A_{ij} \times Educ_i + d_3 Educ_i + \varepsilon_{ij},$$

where $\tilde{A}_{ij}$ is parent i 's baseline beliefs about child j 's academic performance, A_{ij} is child j 's academic performance, and $Educ_i$ is household-average years of parental education. The prediction is $d_2 > 0$: more-educated parents have less attenuated

²⁶Many evaluations use self-reported enrollment as the outcome of interest (e.g., Bourguignon, Ferreira, and Leite 2003; Schultz 2004), but Baird and Özler (2012) show that self-reported and school data do not always match. I have dropout data from 10 percent of the schools and, reassuringly, the coefficient on $Treat \times Score$ is the same regardless of the data source used, reflecting a high correlation between measures (0.5). Since the dependent variable mean is near 1, online Appendix Table C.14 shows robustness to using a probit. Online Appendix Table C.15 shows robustness to including household fixed effects and child-level controls interacted with treatment. For dropouts, we lose statistical power quickly, but the coefficient stays stable and the p -value remains ≤ 0.15 .

TABLE 4—HETEROGENEITY BY PARENT EDUCATION IN BELIEF INACCURACIES AND TREATMENT EFFECTS

	Parent beliefs about child's score in:					
	Overall (1)	Math (2)	English (3)	Chichewa (4)	Math–English (5)	Child 2 – 1 (6)
<i>Panel A. Heterogeneity in belief accuracy</i>						
Score	0.014	0.018	0.014	0.0098	0.013	0.017
× Parent years of education	[0.0038]	[0.0038]	[0.0043]	[0.0036]	[0.0047]	[0.0049]
Score	0.25	0.13	0.22	0.20	0.091	0.32
	[0.023]	[0.023]	[0.026]	[0.021]	[0.029]	[0.028]
Parent years of education	−0.53	−0.98	−0.065	−0.32	−0.78	0.044
	[0.20]	[0.20]	[0.21]	[0.23]	[0.094]	[0.12]
Observations	5,220	5,222	5,222	5,222	5,222	5,218
<i>Panel B. Heterogeneity in the treatment effect on the slope, by parent education, for selected experimental outcomes</i>						
	Endline beliefs (1)		Math workbook difficulty level (2)		English workbook difficulty level (3)	
Treat × Score	−0.026		−0.12		−0.066	
× Parent years of education	[0.0071]		[0.027]		[0.029]	
Treat × Score	0.53		1.92		1.57	
	[0.044]		[0.16]		[0.17]	
Score	0.022		0.079		0.033	
× Parent years of education	[0.0051]		[0.020]		[0.022]	
Score	0.21		0.28		0.61	
	[0.031]		[0.11]		[0.13]	
Treat	1.22		6.54		2.31	
× Parent years of education	[0.39]		[1.45]		[1.53]	
Treat	−31.8		−121.8		−79.2	
	[2.31]		[8.57]		[8.59]	
Parent years of education	−0.79		−3.89		−0.34	
	[0.27]		[1.08]		[1.18]	
Observations	5,208		5,203		5,203	
R^2	0.342		0.221		0.208	
p -val: Treat × Score	0.000		3.7e-06		0.022	
× Years education = 0						
Score used	Overall		Math		English	

Notes: Data sources are baseline survey and baseline test score data. Each observation is a child. Standard errors are clustered at the household level. *Parent years of education* is the household-average years of parental education. Panel A displays regressions of parents' baseline beliefs about their children's true score on their children's true score, average parental years of education, and the interaction. The prediction is that true scores will be more highly correlated with the beliefs of more-educated parents, and that the coefficient on *Score* × *Parent years of education* will be positive. Panel B shows the heterogeneity by parent education in the information treatment effect on the gradient of the investment function. The dependent variable in Column 1 corresponds to the parent's endline beliefs about the child's overall score on a hypothetical test taken the same day as the endline survey. Panel B regressions control for school FE, the between-child score gap, child gender, grade FE, parent gender, and parental education proxy used for stratification.

beliefs. The table shows that d_2 is strongly positive. The magnitudes of the estimates suggest that going from 2 to 7 household-average years of education (the twenty-fifth percentile to the seventy-fifth percentile of the distribution) increases the slopes by roughly 25–55 percent.²⁷ Although less-educated parents have significantly less accurate beliefs than more-educated parents, they are not significantly more

²⁷ An alternate way to look at belief accuracy is to test whether the absolute value of the gap between beliefs and true scores is larger for less-educated parents. Online Appendix Table C.1 presents this test with consistent

overconfident (online Appendix Table C.1, columns 5–6). Note that this result is not inconsistent with estimating a negative d_3 in equation (7): d_3 is not a group mean but rather a group intercept that is mechanically linked with the slope conditional on the mean (i.e., since both $\bar{A}_{ij}$ and A_{ij} are positive, increasing the line's slope without changing its mean decreases its intercept).

RESULT D2: *Information has a larger effect among less-educated parents.*

I now examine whether, in addition to having less accurate beliefs, less-educated parents also respond more to information. I first examine belief updating, shown in panel B of Table 4, column 1, which shows the results of estimating equation (6) fully interacted with household-average years of parent education. Less-educated parents update their beliefs more than more-educated parents. We may also wish to know whether more belief updating translates into “larger effects” on decisions. Testing this is nontrivial, since it is difficult to define exactly what a “larger effect” is. In particular, the magnitude of a parent's response will depend on her preferred investment function, which in turn depends on her preferences and budget, both of which may vary with SES. Some preferred investment functions may be steeper for richer parents, others for poorer parents. To make an “apples to apples” comparison, then, it is useful to focus on the investments where the preferred investment function is as homogeneous as possible across parental education levels, and where the heterogeneity in treatment effects thus speaks directly to heterogeneity in the impact of information, since, *if there were no gaps* in information, there would be no heterogeneity in the results by parental education. The choice of difficulty level of free workbooks is most likely to meet this criterion, and was expressly included in the design to provide homogeneity across education levels.²⁸ Columns 2 and 3 display the results, showing that information has a larger effect for less-educated parents. At baseline, the workbook choices of above-median-education parents are roughly 90 percent (30 percent) more steeply sloped for math (English) than the choices of below-median-education parents (see positive coefficient on $Score \times Parent\ years\ of\ education$); information fully closes the gap (see negative coefficient on $Treat \times Score \times Parent\ years\ of\ education$). See Dizon-Ross (2018) to see the results for the other outcomes, where there is more potential for heterogeneity in the preferred investment function by parental education.

E. Welfare and Average Treatment Effects

This experiment shows that providing information to parents affects their beliefs and decisions. This reveals that parents' decisions at baseline did not fully incorporate the information, and that parents themselves think the information makes them better off, i.e., that, according to their own (perceived) utility functions, information

results. Online Appendix Tables C.4 and C.16 show robustness to using other measures of parent education and child performance and to controlling for other variables (including school fixed effects) interacted with score.

²⁸ Since the workbooks are free, wealth should not affect the choice. Moreover, we expect parents to choose the workbook most closely matched to their beliefs about their child's performance, and there is no reason to expect that to vary by parental education. This is corroborated by online Appendix Table C.17: control group regressions of workbook difficulty level on baseline beliefs show no heterogeneity by parent education.

increases utility. Although establishing these findings was the key conceptual goal of the paper, one may also wonder about the broader welfare implications. Welfare conclusions are difficult, since any intervention that corrects one market imperfection can decrease welfare if there are multiple interacting market failures (the “theory of the second best”). Definitive welfare conclusions are thus beyond the scope of this paper. That said, I now provide some speculative discussion.

First, to draw conclusions about *parents’* welfare, we would need to know whether their perceived education production function is correct, including how predictive the performance information is for the returns to investment.²⁹ Reassuringly, in both developed and developing countries, grades are correlated with other measures of ability and may even be better predictors of life outcomes (e.g., wages) than standard measures of ability (Borghans et al. 2011, 2016; Sternberg et al. 2001). In addition, for the outcomes that proxy for years of schooling (primary school enrollment and the secondary school lottery), although there are no estimates of the production function in Malawi, estimates from other contexts suggest that years of schooling and other measures of ability are complements (Pitt, Rosenzweig, and Hassan 1990; Aizer and Cunha 2012). There are also reasons to expect that the complementarity might be greater in this setting: the education system in sub-Saharan Africa is particularly oriented toward high performers (Duflo, Dupas, and Kremer 2011), and achievement tests determine progression through school and access to higher levels of schooling, thus mechanically linking them with returns.³⁰ My finding that parents allocate more years of schooling to their higher performers suggests they believe this complementarity exists and is therefore *consistent* with parents being correct about the production function, but more evidence would be needed for firm conclusions.

Second, to move to social welfare, we would also need to assess the externalities of education and the social welfare function. Some students receive higher investments as a result of the intervention and some lower; if child welfare were increasing in schooling, then the social welfare impact would depend on the welfare weights placed on those children. More broadly, although no single summary statistic can fully capture the welfare impacts, the conventional wisdom is that collectively market imperfections (such as positive externalities of education) cause the average level of education to be below the optimum. As a result, it might be particularly concerning if providing information about academic performance caused the average level of investments to fall. However, reassuringly, information does not decrease the average level of investments, although the results are imprecise. Panel A of Appendix Table A.1 shows that there is no statistically significant ATE of information on the investments that could proxy for the overall level of investments: enrollment, expenditures, and attendance.³¹

One might be surprised by the absence of an ATE for enrollment. On average, parents overestimate their children at baseline and, for enrollment, invest more in

²⁹ Note that this is much less of a concern when analyzing some of the experimental outcomes (e.g., the workbooks and remedial textbooks). These investments were designed to have clear predictions for increased returns, and, across the parental education spectrum, parents’ reallocations align with the predictions.

³⁰ That said, even in completely different educational markets characterized by heavy penetration of private schools, schools also appear to cater to richer and higher-achieving students; thus, the level of complementarity may be similar in those types of markets as well (Bau 2017).

³¹ The workbooks and textbooks were specifically designed to look at the *types* of investment chosen and thus their level does not proxy for overall spending. For completeness, however, these ATEs are also reported.

higher performers, suggesting that information might decrease enrollment. Online Appendix K describes several potential explanations for the lack of an ATE, one of which I find empirical support for: parents respond more to information when the information is positive than when it is negative. This is consistent with the findings of the motivated beliefs literature (Eil and Rao 2011, Mobius and Rosenblat 2014). See online Appendix K and online Appendix Table C.18 for analysis and discussion.

F. Uncertainty

Does information affect investments primarily by affecting the mean or the uncertainty of parents' beliefs distributions? I investigate this in online Appendix L. The analyses suggest that the primary mechanism for information's effects on the *types* of investments chosen (e.g., difficulty levels of workbooks) is changes to the mean/accuracy of beliefs, but that changes in the uncertainty of beliefs also play a role for the larger investments that proxy more for the *level* of investment.

IV. Conclusion

This paper highlights an important source of misinformation that affects decisions: parents' inaccurate beliefs about their children's academic performance. I show that (perceived) academic performance is an important input into parents' investment decisions, but that despite the ready availability of academic performance information, many parents' beliefs about their children's academic performance are quite wrong, with important consequences for the allocation of educational investments. Providing academic performance information to parents causes them to change both the *level* and *type* of investments they choose for their children. The impacts are seen across a broad range of investments, from those with very clean predictions about how parents should invest to maximize returns (e.g., remedial textbooks that are more useful for low-performing students), to more consequential investments that proxy for overall educational attainment.

It is perhaps surprising that baseline information is poor if the returns to knowledge are high and the information is, in principle, readily available. But parents may overestimate their own knowledge, or the (perceived) costs of acquiring information may be high, especially for uneducated or illiterate parents. Indeed, interviews with parents suggest that uneducated parents are intimidated to talk with their children's teachers. These barriers to accessing information may be more pervasive in poorer countries and among poorer households. Consistent with this, I find that less-educated parents in my setting have less accurate beliefs. This same pattern is also evident in beliefs data from other contexts, including the United States,³² and suggests a potential link between poor access to information and poverty.

This paper focuses on identifying the causal chain between parents' beliefs and their investments. One area for future research would be to extend the causal chain further to better understand the link between investments and welfare, for example, by measuring the objective returns to different educational investments

³² US data were provided by Alexander and Entwisle (2006) and analyzed by the author.

and comparing them with parents' perceptions. A second area would be to evaluate the impact of information on long-run educational outcomes, both on average and across the distribution of performance. In settings like this, where parents believe years of schooling and performance are complements, information should unambiguously increase outcomes for high-performers. However, the effects on low-performers are ambiguous: information may decrease the level of investments but allow parents to more efficiently choose the right *types* of investments. Understanding these impacts is relevant for policy: if educational outcomes fall for poorly-performing students, that may be optimal for the families who prefer to decrease investments in their poorly-performing child, but policymakers often want to improve educational outcomes for poorly performing students. Finally, a third direction for future work would be to design and test a scalable information-dissemination strategy.

APPENDIX

TABLE A.1—AVERAGE TREATMENT EFFECTS

	Experimental outcomes				Non-experimental outcomes		
	Endline beliefs (1)	Math workbook difficulty level (2)	English workbook difficulty level (3)	ln(math textbook WTP) – ln(English textbook WTP) (4)	Enrollment (5)	ln(Total educ. expenditures) (6)	Attendance rate (7)
<i>Panel A. Average treatment effects</i>							
Treat	–7.47 [0.53]	–32.4 [2.16]	–13.7 [2.25]	0.14 [0.041]	–0.39 [0.71]	0.0024 [0.049]	–0.16 [0.79]
Observations	5,244	5,239	5,239	5,219	1,786	1,709	1,827
<i>Panel B. Uncertainty level effects: beliefs within 10 points of truth</i>							
Treat	–0.53 [0.91]	–10.3 [4.05]	1.59 [3.67]	0.067 [0.072]	–0.47 [0.89]	0.070 [0.086]	0.84 [1.32]
Observations	1,572	1,299	1,657	1,589	579	550	541

Notes: Data sources are baseline survey, baseline test score data, both endline surveys, and endline administrative data. Each observation is a child. Standard errors are clustered at the household level. Regressions control for school FE, average parental years of education, parent gender, a parental education proxy used for stratification, the between-child score gap, child baseline performance, child gender, grade FE, and the baseline value of the dependent variable (baseline value not available for enrollment or experimental outcomes). The dependent variable in Column 1 corresponds to the parent's endline beliefs about the child's overall score on a hypothetical test taken the same day as the endline survey. Workbook difficulty choices are coded as 0 for beginner, 100 for average, 200 for advanced. *Enrollment* is defined as being enrolled in school one year after the intervention; enrollment and attendance scaled to be out of 100 (so, for example, enrollment is equal to 100 if the child is still enrolled and 0 otherwise). Panel B uses the relevant measure of beliefs (e.g., overall for beliefs, math – English for textbooks; see Table 3 for details).

A. Discussion and Proofs of Predictions 1 and 2

I begin by proving and discussing the predictions in the case that the preferred investment function is linear: $s^*(\alpha) = \beta_0 + \beta_1 \alpha$. I then move on to the nonlinear case.

Linear Case

Prediction 1: Attenuation in the slope of the actual investment function.—We want to show that if beliefs are an “attenuated” function of true performance (i.e., have a slope less than 1 if regressed on true performance), then the slope of the actual investment function will be lower in magnitude than the slope of the preferred function. The slope of the preferred investment function (i.e., the slope of investments on beliefs) is β_1 . Using the standard OLS formula, the slope of the actual investment function (i.e., the slope of expected investments on true performance) is $\text{cov}(\beta_0 + \beta_1 \alpha, a) / \text{var}(a) = \beta_1 (\text{cov}(\alpha, a) / \text{var}(a))$. Thus, whenever $\text{cov}(\alpha, a) / \text{var}(a) \neq 1$, inaccurate beliefs will cause the actual slope to differ from the preferred slope, and whenever $\text{cov}(\alpha, a) / \text{var}(a) < 1$, there is attenuation. Since $\text{cov}(\alpha, a) / \text{var}(a)$ is the slope from regressing believed performance on true performance, this means that the condition for attenuation in the slope of $\tilde{s}(\cdot)$ is that beliefs are an attenuated function of true performance, i.e., have a slope less than 1.

Prediction 2: If there is baseline attenuation, information increases the slope of investments.—In the linear case, chosen investments equal $s^*(\alpha) + \varepsilon = \beta_0 + \beta_1 \alpha + \varepsilon$. I first outline the bias in an observational data approach, and then outline how an experiment addresses this bias. The observational approach would be to compare the slopes estimated from regressing baseline (or control group) s on α with the slope from regressing baseline s on a . The slope from regressing on α will be the true causal slope, β_1 , plus an omitted variable bias (OVB) term, $\text{cov}(\alpha, \varepsilon) / \text{var}(\alpha)$. The slope from regressing on a will be the true causal slope derived above, $\beta_1 (\text{cov}(\alpha, a) / \text{var}(a))$, plus an OVB term: $\text{cov}(a, \varepsilon) / \text{var}(a)$. Thus, the difference in slopes will be $(\beta_1 - \beta_1 (\text{cov}(\alpha, a) / \text{var}(a))) + ((\text{cov}(\alpha, \varepsilon) / \text{var}(\alpha)) - (\text{cov}(a, \varepsilon) / \text{var}(a)))$ and so will only give us an unbiased estimate of the true difference in slopes, $\beta_1 - \beta_1 (\text{cov}(\alpha, a) / \text{var}(a))$, if the second term (i.e., the difference between the OVB terms $((\text{cov}(\alpha, \varepsilon) / \text{var}(\alpha)) - (\text{cov}(a, \varepsilon) / \text{var}(a)))$) is equal to 0.

An experiment can solve this problem. Consider comparing the slopes of the actual investment functions (s regressed on a) for parents who have received information about a (treatment group) versus those who have not (control group). Parents in the treatment group will now base investments on true performance a , so their investments will be $s^*(a) + \varepsilon = \beta_0 + \beta_1 a + \varepsilon$.³³ The slope in the treatment group will thus be $\beta_1 + (\text{cov}(a, \varepsilon) / \text{var}(a))$, whereas in the control group, it will be the same as above: $\beta_1 (\text{cov}(\alpha, a) / \text{var}(a)) + (\text{cov}(a, \varepsilon) / \text{var}(a))$. Since, unlike for the observational approach, the omitted variable terms are now identical, comparing the slope between treatment and control groups will allow us to estimate the true difference in slopes $|\beta_1 - \beta_1 (\text{cov}(\alpha, a) / \text{var}(a))|$. If investments were attenuated at

³³Note that this assumes that parents fully update their beliefs in response to the intervention. If they only partially update their beliefs, then the difference in slope between treatment and control groups would be weighted downward by the updating parameter (i.e., if updated beliefs were a weighted combination of a and α with γ the weight on a , then the difference in slopes would uncover $\gamma(\beta_1 - \beta_1 (\text{cov}(\alpha, a) / \text{var}(a)))$).

baseline, that difference will be positive, meaning that information will increase the magnitude of the slope.

Nonlinear Case

I now outline corollaries for Prediction 1 and for Prediction 2 that outline the tests to perform when the investment function is nonlinear. I make two assumptions. First, I assume that investments are monotonic and never flat in performance; monotonic implies that investments do not switch between being a complement with performance in some regions and a substitute in others. Second, I assume that the true and believed performance distributions have the same support.

When beliefs are attenuated, if they affect investment, then many parents at the maximum of the performance distribution are choosing investments more appropriate for lower-performing children, and vice versa at the minimum. As a result, the absolute value of the gap in investments between children at the top and bottom of the performance distribution is compressed relative to the preferred gap. In the linear case, this implies that the OLS slope of expected investments will be flatter on true than believed performance (Prediction 1). However, in the nonlinear case, if beliefs are attenuated, the OLS slope may *not* be flatter on true than believed investment even if beliefs affect investment.³⁴ It *may* be flatter, in which case that attenuation would provide evidence that belief inaccuracies affect investments; however, if it is not, we can test directly for compression at the ends of the distribution as an alternate test for whether beliefs affect investments.

Since we are focused on the ends of the distribution, the relevant definition for “attenuation” of beliefs now also focuses on the ends of the distribution. Define $\bar{a}$, and $\underline{a}$ as the maximum and minimum of the true performance distribution. We now have the following testable analog of Prediction 1 for the case where the investment function is nonlinear.

COROLLARY 1: *Assume that there is “attenuation” of beliefs at the ends of the performance distribution, i.e., that $E[\alpha|a = \bar{a}] < \bar{a}$ and $E[\alpha|a = \underline{a}] > \underline{a}$. Then, if investments depend on beliefs, the gap in investments between children at the top and bottom of the true performance distribution is compressed relative to the gap between children at the top and bottom of the believed performance distribution: $|E[s|a = \bar{a}] - E[s|a = \underline{a}]| < |E[s|\alpha = \bar{a}] - E[s|\alpha = \underline{a}]|$.*

PROOF:

Without loss of generality, I describe the case where investments are increasing in performance. The maximum chosen investment across the population is $s^*(\bar{a})$, and, for any belief $\hat{\alpha} < \bar{a}$, $s^*(\hat{\alpha}) < s^*(\bar{a})$. Since $E[\alpha|a = \bar{a}] < \bar{a}$, some parents at the top of the performance distribution (i.e., for whom $a = \bar{a}$) have beliefs that are not at the top of the distribution ($\alpha < \bar{a}$) and thus are choosing investments less than

³⁴ Whether the OLS slope will be flatter depends on the relative density of believed and true performance in the areas where the slope of the s^* function is steeper.

$s^*(\bar{a})$. As a result $E[s|a = \bar{a}] < E[s|\alpha = \bar{a}] = s^*(\bar{a})$. One can use an analogous argument to show that $E[s|a = \underline{a}] > E[s|\alpha = \underline{a}] = s^*(\underline{a})$. ■

In the nonlinear case, by the same logic underlying Prediction 2, providing information will also undo the attenuation or compression at the ends of the distribution, yielding the following testable analog of Prediction 2.

COROLLARY 2: *If (i) the baseline gap in investments between children at the top and bottom of the true performance distribution is compressed due to inaccurate beliefs and (ii) information causes parents to update their beliefs, then providing information on true performance will increase the gap in investments between children at the top and bottom of the believed performance distribution, i.e., increase $|E[s|a = \bar{a}] - E[s|a = \underline{a}]|$.*

This follows because the parents at the top and bottom of the performance distribution will now *believe* that they are at the top and bottom of the performance distribution and choose investments accordingly. Hence, after parents receive information, $E[s|a = \bar{a}] = E[s|\alpha = \bar{a}]$ and $E[s|a = \underline{a}] = E[s|\alpha = \underline{a}]$.

B. Uncertainty Predictions

There are many ways to model uncertainty in beliefs. Here, I show one potential framework which yields the prediction that higher uncertainty in parents' beliefs about academic performance, σ^2 , leads to greater attenuation in the slope of investments on mean beliefs, $|\partial s^{**}/\partial \alpha|$. The framework captures the intuition described in the main text: that uncertainty may make parents hesitate to rely as strongly on their mean beliefs when making their investment decisions. This is a richer model than the one used in Section I.

Assume there is some true unobserved underlying academic potential. Call this $\mathbf{a}$ and call parents' beliefs about it α . Assume this underlying academic potential is what determines returns and is thus what parents truly want to base decisions on. Assume further that academic potential is distinct from academic performance, a , where a is what we measured baseline beliefs on, and what we delivered information about in the intervention; instead, academic performance a is taken by parents as a signal of $\mathbf{a}$.

In this context, we can model beliefs about academic potential α as being a convex combination of beliefs about school performance, α , and beliefs about all other aspects or signals of academic potential, $\alpha_{-\alpha}$, given by

$$\alpha = \lambda \alpha + (1 - \lambda) \alpha_{-\alpha},$$

where λ is the weight on the academic performance.

Since preferred investments would be a function of α , not α , we could write the preferred investment function as $\tilde{s}^*(\alpha)$. For expositional simplicity, let's look at the linear case where $\tilde{s}^*(\alpha) = \beta_0 + \beta_1 \alpha$ (where the β_1 notation distinguishes this from the preferred investment function in the simpler model from Section I and

Appendix Section A) and where uncertainty in α itself does not affect investments (only uncertainty in α affects investments).³⁵ Preferred investments could then be written as

$$\begin{aligned}\tilde{s}^*(\alpha) &= \tilde{\beta}_0 + \tilde{\beta}_1\alpha \\ &= \tilde{\beta}_0 + \tilde{\beta}_1\lambda\alpha + \tilde{\beta}_1(1-\lambda)\alpha_{-\alpha}.\end{aligned}$$

In this context, providing information about academic performance, a , should decrease σ^2 , i.e., increase the certainty of parents' beliefs about a , as represented through the beliefs distribution $g(\alpha, \sigma^2)$. This could increase the weight that parents place on beliefs about academic performance α when forming their beliefs about underlying academic potential α , that is, increase λ . Since λ increases, under most assumptions for the form that $\alpha_{-\alpha}$ would take,³⁶ the slope of investments on beliefs about school performance α should also increase.

Note that this is a channel for uncertainty to change the slope of investments on beliefs about academic performance, α , even if the underlying slope of the true preferred investment function on beliefs about academic potential, $\tilde{a}$, does not change.³⁷

REFERENCES

- Aizer, Anna, and Flávio Cunha. 2012. "The Production of Child Human Capital: Endowments, Investments and Fertility." Unpublished.
- Alexander, Karl L., and Doris R. Entwisle. 2006. "Baltimore Beginning School Study, 1982–2002." Harvard-MIT Data Centers. Henry A. Murray Research Archive. Log#01293.
- Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja. 2016. "Report Cards: The Impact of Providing School and Child Test Scores on Educational Markets." Unpublished.
- Arcidiacono, Peter, V. Joseph Hotz, and Songman Kang. 2012. "Modeling College Major Choices Using Elicited Measures of Expectations and Counterfactuals." *Journal of Econometrics* 166 (1): 3–16.
- Baird, Sarah, and Berk Özler. 2012. "Examining the Reliability of Self-Reported Data on School Participation." *Journal of Development Economics* 98 (1): 89–93.
- Banerjee, Abhijit V., Rukmini Banerji, Esther Duflo, Rachel Glennerster, and Stuti Khemani. 2010. "Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in Education in India." *American Economic Journal: Economic Policy* 2 (1): 1–30.
- Bau, Natalie. 2017. "School Competition and Product Differentiation." Unpublished.
- Bergman, Peter. 2016. "Parent-Child Information Frictions and Human Capital Investment: Evidence from a Field Experiment." Unpublished.
- Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and Katherine L. Milkman. 2015. "The Effect of Providing Peer Information on Retirement Savings Decisions." *Journal of Finance* 70 (3): 1161–201.
- Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2012. "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block FAFSA Experiment." *Quarterly Journal of Economics* 127 (3): 1205–42.

³⁵In the notation of Section I, $\tilde{s}^{**}(\alpha, \text{var}(\alpha)) = \tilde{s}^*(\alpha)$. For example, parents' maximand when choosing s given academic potential $\mathbf{a}$ could be quadratic loss in $\mathbf{a}$: $U(f(s_{ij}, \mathbf{a}_{ij})) = -c(\gamma s_{ij} - \mathbf{a}_{ij})^2$.

³⁶Specifically, the regression of investments on α would have slope $\lambda\tilde{\beta}_1 + (1-\lambda)\tilde{\beta}_1(\text{cov}(\alpha, \alpha_{-\alpha})/\text{var}(\alpha)) = \lambda\tilde{\beta}_1 + (1-\lambda)\tilde{\beta}_1\text{corr}(\alpha, \alpha_{-\alpha})(\text{std}(\alpha_{-\alpha})/\text{std}(\alpha))$. Thus, since $\text{corr}(\alpha, \alpha_{-\alpha}) \leq 1$, increasing λ should increase the slope as long as the variance of $\alpha_{-\alpha}$ is not too much larger than the variance of α .

³⁷It is useful to note that in this richer model, although providing information about a should unambiguously increase the certainty of α , it is ambiguous whether it will decrease or increase the uncertainty of beliefs about $\mathbf{a}$. For example, if the information were very different from parents' prior beliefs, it could increase the uncertainty of beliefs about $\mathbf{a}$.

- Bhargava, Saurabh, and Dayanand Manoli.** 2015. "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105 (11): 3489–529.
- Bobba, Matteo, and Verónica Frisancho.** 2016. "Learning about Oneself: The Effects of Performance Feedback on School Choice." Unpublished.
- Borghans, Lex, Bart H. H. Golsteyn, James Heckman, and John Eric Humphries.** 2011. "Identification Problems in Personality Psychology." *Personality and Individual Differences* 51 (3): 315–20.
- Borghans, Lex, Bart H. H. Golsteyn, James J. Heckman, and John Eric Humphries.** 2016. "What Grades and Achievement Tests Measure." *Proceedings of the National Academy of Sciences* 113 (47): 13354–9.
- Bourguignon, François, Francisco H. G. Ferreira, and Phillippe G. Leite.** 2003. "Conditional Cash Transfers, Schooling, and Child Labor: Micro-Simulating Brazil's Bolsa Escola Program." *World Bank Economic Review* 17 (2): 229–54.
- Chetty, Raj, and Emmanuel Saez.** 2009. "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients." NBER Working Paper 14836.
- Chevalier, Arnaud, Steve Gibbons, Andy Thorpe, Martin Snell, and Sherria Hoskins.** 2009. "Students' Academic Self-Perception." *Economics of Education Review* 28 (6): 716–27.
- de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth.** 2017. "Measuring and Bounding Experimenter Demand." NBER Working Paper 23470.
- Dinkelmann, Taryn, and Claudia 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): 244–57.
- Dizon-Ross, Rebecca.** 2018. "Identifying the (Perceived) Education Production Function: Evidence from Information Shocks." Unpublished.
- Dizon-Ross, Rebecca.** 2019. "Parents' Beliefs about Their Children's Academic Ability: Implications for Educational Investments: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20171172>.
- Duflo, Esther, Pascaline Dupas, and Michael Kremer.** 2011. "Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya." *American Economic Review* 101 (5): 1739–74.
- Duflo, Esther, and Emmanuel Saez.** 2002. "Participation and Investment Decisions in a Retirement Plan: The Influence of Colleagues' Choices." *Journal of Public Economics* 85 (1): 121–28.
- Dupas, Pascaline.** 2011. "Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics* 3 (1): 1–34.
- Eil, David, and Justin M. Rao.** 2011. "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself." *American Economic Journal: Microeconomics* 3 (2): 114–38.
- Fitzsimons, Emla, Bansi Malde, Alice Mesnard, and Marcos Vera-Hernández.** 2016. "Nutrition, Information and Household Behavior: Experimental Evidence from Malawi." *Journal of Development Economics* 122: 113–26.
- Hoxby, Caroline, and Sarah Turner.** 2013. "Expanding College Opportunities for High-Achieving, Low-Income Students." Stanford Institute for Economic Policy Research Discussion Paper 12–014.
- Jamison, Julian C., Dean Karlan, and Pia Raffler.** 2013. "Mixed Method Evaluation of a Passive mHealth Sexual Information Texting Service in Uganda." NBER Working Paper 19107.
- Jensen, Robert.** 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics* 125 (2): 515–48.
- Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian Wrobel.** 2011. "Comparison Friction: Experimental Evidence from Medicare Drug Plans." NBER Working Paper 17410.
- Liebman, Jeffrey B., and Erzo F. P. Luttmer.** 2015. "Would People Behave Differently If They Better Understood Social Security? Evidence from a Field Experiment." *American Economic Journal: Economic Policy* 7 (1): 275–99.
- Mobius, Markus, and Tanya Rosenblat.** 2014. "Social Learning in Economics." *Annual Review of Economics* 6: 827–47.
- Nguyen, Trang.** 2008. "Information, Role Models and Perceived Returns to Education: Experimental Evidence from Madagascar." Unpublished.
- Paviot, Laura, Nina Heinsohn, and Julia Korkman.** 2008. "Extra Tuition in Southern and Eastern Africa: Coverage, Growth, and Linkages with Pupil Achievement." *International Journal of Educational Development* 28 (2): 149–60.

- Pitt, Mark M., Mark R. Rosenzweig, and Md. Nazmul Hassan.** 1990. "Productivity, Health, and Inequality in the Intrahousehold Distribution of Food in Low-Income Countries." *American Economic Review* 80 (5): 1139–56.
- Rogers, Todd, and Avi Feller.** 2016. "Intervening through Influential Third Parties: Reducing Student Absences at Scale via Parents." Unpublished.
- Schultz, T. Paul.** 2004. "School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program." *Journal of Development Economics* 74 (1): 199–250.
- Sternberg, Robert J., Catherine Nokes, P. Wenzel Geissler, Ruth Prince, Frederick Okatcha, Donald A. Bundy, and Elena L. Grigorenko.** 2001. "The Relationship between Academic and Practical Intelligence: A Case Study in Kenya." *Intelligence* 29 (5): 401–18.
- Stinebrickner, Ralph, and Todd R. Stinebrickner.** 2014. "A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout." *Review of Economic Studies* 81 (1): 426–72.
- Stinebrickner, Todd, and Ralph Stinebrickner.** 2012. "Learning about Academic Ability and the College Dropout Decision." *Journal of Labor Economics* 30 (4): 707–48.
- UNESCO.** 2013. "Education: Mean Years of Schooling." <http://data.uis.unesco.org/index.aspx?queryid=242> (accessed November 1, 2014).
- Wiswall, Matthew, and Basit Zafar.** 2015. "Determinants of College Major Choice: Identification Using an Information Experiment." *Review of Economic Studies* 82 (2): 791–824.
- World Bank.** 1998. *World Development Report 1998–1999: Knowledge for Development*. New York: Oxford University Press.
- World Bank.** 2008. *World Development Report 2008: Agriculture for Development*. Washington, DC: World Bank.
- World Bank.** 2010. "The Education System in Malawi." Washington, DC: World Bank.