When Can School Inputs Improve Test Scores?

Jishnu Das (DECRG, The World Bank)
Stefan Dercon (Oxford University)
James Habyarimana (Harvard University)
Pramila Krishnan (Cambridge University)*

October 16, 2003

Abstract

The relationship between school inputs and educational outcomes is critical for educational policy. We recognize that households will respond optimally to changes in school inputs and study how such responses affect the link between school inputs and cognitive achievement. To incorporate the forward-looking behavior of households, we present a household optimization model relating household resources and congnitive achievement to school inputs. In this framework if household and school inputs are technical substitutes in the production function for cognitive achievement, the impact of unanticipated inputs is larger than that of anticipated inputs. We test the predictions of the model for *non-salary cash grants to schools* using a unique dataset from Zambia. We find that household educational expenditures and school cash grants are substitutes with a coefficient of elasticity between -0.27 and -0.5. Consistent with the optimization model, anticipated funds have no impact on cognitive achievement, but unanticipated funds lead to significant improvements in learning. This methodology has important implications for educational research and policy.

Draft: Preliminary and Incomplete

*Corresponding Author: Jishnu Das (jdas1@worldbank.org). We thank Hanan Jacoby and workshop participants at the World Bank for comments. Funding for the survey was provided by DFID.

1. Introduction

In the last two decades the relationship between schooling inputs and educational outcomes has received considerable attention in academic and policy forums. Although it is recognized that households play a critical role in the determination of such outcomes, the literature has bifurcated in two distinct strands.[1] One strand, concerned with the impact of school inputs on cognitive achievement, has focused on estimating educational production functions where cognitive achievement is determined as a function of schooling inputs. The second strand is concerned with the effect of household characteristics on cognitive achievement *independently* of school inputs. Curiously, the notion that household responses themselves affect the relationship between cognitive achievement and school inputs has received little attention in the literature.

As examples of how this may happen, consider the following quotes:

"This year the school did not provide any textbooks and we had to buy our own" (Household Interview: Zambia)

"I remember when the school changed its curriculum, they decreased their emphasis on science. I offered to teach (and taught) a special class for children who were interested" (Personal Communication: US)

Both quotes illustrate a simple point: Households take into account resources at the school level and plan their actions accordingly. The first quote suggests that household responses will attenuate the relationship between textbook provision and cognitive achievement; the second does the same for broader changes regarding the school curriculum. In a recent paper Todd and Wolpin (2003) point out that in the presence of household responses, estimates based on the production function approach will capture a 'policy-effect' that incorporates both the marginal impact of school inputs on outcomes *as well as* household responses to such inputs. This raises an important question that defines the central problem addressed here: How are we to understand the relationship between school inputs and cognitive achievement in an environment where households respond to the provision of such inputs?

To examine this question we first derive a household optimization model of cognitive achievement. The model incorporates the basic assumption that households respond optimally to the provision of inputs at the school level. However, we also need to incorporate the notion that households may be *forward-looking* so that responses occur not only in the period that school inputs are provided, but the moment that new information becomes available. Our explicit consideration of the dynamics has important implications—since household adjustments occur with any new information, production function parameters can be identified only through the impact of unanticipated inputs on cognitive achievement.

In this framework, we show that the impact of school inputs depends on (a) whether they are anticipated or not and (b) the extent of substitutability between household and school inputs in the production function for cognitive achievement. If household and school inputs are technical substitutes, an anticipated increase in inputs in the next period increases household contributions in the current period and decreases them in

[1] The path-breaking Coleman Report (United States National Center for Educational Statistics, 1966) for instance stressed the importance of household characteristics for child achievement.

the next, whereas if they are technical complements, the impact of anticipated increases in school inputs on current contributions depends on the strength of the households' preferences for a smooth consumption path. Unanticipated increases in school inputs in the next period preclude household responses in either the current period or the next. These differences lead to a testable prediction: If household and school inputs are (technical) substitutes, unanticipated inputs will have a larger impact on cognitive achievement than anticipated inputs; if they are complements, the relative effects depend on household preferences.

We test this prediction using data from Zambia in 2002-2003. The educational environment in the country is particularly well suited for our empirical exercise. The system is largely based on public schools (less than 2% of all schools are privately run) and the country has a history of high enrollment rates and school participation, suggesting that household involvement in children's education is high.[2] In 2000 the government legislated a fixed cash grant to every school. These grants were large: After accounting for household responses, they represented 150% of household level educational expenditures for the lowest wealth deciles and 19% for the top wealth decile. Moreover the simplicity of the allocation rule ensured that the grants reached their intended recipients (see Das, Dercon, Habyarimana and Krishnan, 2003), suggesting that in the year of the survey the fixed cash grants would be anticipated by households making their educational investment decisions for the year.

In addition schools could also receive cash from other sources, but these alternative sources were highly unreliable and unpredictable. In the year of the survey, less than 25% of all schools received any such grants and conditional on receipt, there was tremendous variation with some schools receiving 30 times as much as others. Apart from cash, few resources were distributed to schools during the year of the survey. Finally, following an agenda of 'free' education, all institutionalized parental contributions to schools were banned in April 2001 (typically these took the form of Parent-Teacher Association fees). Taken together this implied that educational expenditures for children could be met either through cash grants to schools or through direct parental contribution at the level of the household. These two factors present exactly the framework required to test our model with high parental contributions towards educational inputs on the one hand and two different streams of cash disbursements to schools, one steady and the other unpredictable, on the other.

To exploit the characteristics of the educational environment, we collected a unique data set for a representative sample of schools in four provinces of the country (covering 58% of the population). The survey includes data on school inputs as well as two test scores on the same sample of students one year apart. To supplement this data, we also collected information for households matched to a sub-sample of schools identified as 'remote' using GIS mapping tools. This allows us to directly relate household and school inputs in an environment where issues of school choice are eliminated. We are then interested in the effect of anticipated and unanticipated cash grants to schools on household educational expenditures and cognitive achievement.

We find strong support for the household approach to cognitive achievement. Using the matched school-household data, our results suggest that household educational expenditures and cash grants to schools are

[2] Net enrollments are upwards of 80% for both boys and girls (see Figure 2a and 2b).

substitutes. The elasticity of substitution between the two is high and significant, with estimates ranging from –0.45 to -0.55 depending on the specification used. In line with the predictions of our model we then find evidence that unanticipated grants have a significant and substantial impact on the growth of cognitive achievement while the effect of anticipated grants is small and insignificant. These results hold for the subjects of Mathematics and English (although the difference is more pronounced for the latter), and are robust to potential mis-specification arising from omitted variables in the regression.

The significance of our results goes beyond the particular policy environment considered here. A failure to reject the null hypothesis in studies that use the production function approach could arise either because the effect of school inputs on cognitive achievement through the production function is zero *or* because households substitute their own resources for such inputs. While in our case the substitution may take the form of textbooks or writing materials, in a more general setting it may include parental time, private tuition and other inputs. Our results show that the policy-effect of school inputs are different from the production function parameters, and this has important consequences both for estimation techniques and for educational policy, a detailed discussion of which we postpone till later.

This work is new and innovative for a number of reasons. First, the methodology adopted here extends the work of Becker and Tomes (1986) to the determination of cognitive achievement and thus allows us to directly incorporate household responses and school inputs in a single conceptual framework. Second, the unique data collected on matched schools and households permits the direct estimation of household responses to school inputs; while this is clearly an important issue for policy, ours is the one of the first papers to provide an estimate of this relationship in the context of education. Third, the combination of funding patterns in the country and panel data on cognitive achievement provides an excellent opportunity to separate policy effects and production function parameters of schooling inputs; doing so yields new insights on the process through which school inputs may affect educational outcomes. Thus, the combination of the methodology and the unique data collected allows us to provide a firm microeconomic foundation for the relationship between school inputs and cognitive achievement in the context of a household optimization model.

The remainder of the paper is structured as follows. We briefly review the literature in Section II. In section III we present the model and we derive the empirical specification in section IV. Section V presents the results from the matched school-household data and section VI presents the results of anticipated and unanticipated cash grants on the growth of cognitive achievement. Section VII discuss the policy relevance of our work and possibilities for future research.

2. Review of Literature

There are three strands of the literature that we relate to. The first examines the relationship between schooling inputs and cognitive achievement in the context of production functions, the second the impact of household characteristics on educational outcomes. Finally, a third literature has examined the impact of public subsidies on private outcomes, mostly in the context of labor supply and private transfers. We describe each of these briefly.

The literature on educational production functions (for a review, see Hanushek, 1997) attempts to estimate the effect of school inputs on cognitive achievement. The main estimation concern that studies have dealt with is the presence of unobserved heterogeneity, which could contaminate estimates if correlated with the provision of inputs. In response to the omitted variable problem that such heterogeneity creates, studies have tried to exploit 'natural-experiments' (Angrist and Lavy (1999), Case and Deaton (1999), Urquiola (2003)), 'value-added' approaches (Hanushek, 1971) or more recently, randomized treatment-control designs (Banerjee, Cole, Duflo and Linden (2003), Glewwe (2002) provides a review) to argue for causality. Below, we show that our approach allows for greater flexibility in the treatment of unobserved heterogeneity—heuristically, since such heterogeneity is already known at time period $t-1$, it has no effect on growth rates *between* $t-1$ and t.[3]

Further, the methodology adopted here contextualizes results obtained in the educational production function literature in the presence of household responses. Specifically, the policy-effect that is captured in these studies tells us little about *why* certain inputs are successful (or not) in improving cognitive achievement. Clearly this information is important for policy—if the provision of textbooks does not improve test-scores, is it because textbooks are insignificant in the production function or is it due to optimal compensating responses at either the level of the household/school?[4] The household optimization framework makes clear that unanticipated inputs provide the key to understanding these differences.[5]

At the level of the household, studies have examined the relationship between educational outcomes and household characteristics in the context of a household optimization model (Glewwe (2002) and Jacoby and Skoufias (1997)).[6] Fewer studies have examined the role of household characteristics on cognitive achievement; exceptions include Brown (2003), Case and Deaton (1999) and Glewwe and Jacoby (1994) who look at the relationship between parental education and child learning, and Alderman et al.(1997) who examine the effect of household income on test scores. To our knowledge however, there are no studies that examine the role of households in determining the link between school inputs and educational outcomes.

Our findings on household responses do however have an established precedent in the literature on private responses to public transfer programs. The research on this front has typically examined labor supply responses (Moffitt, 1992, Ravallion and Datt, 1995) and private transfers (Cox and Jimenez, 1995 and Rosenzweig and Wolpin, 1994) to find that the effect of government subsidies is generally attenuated through the presence of household responses. The Euler framework developed here has also been used to assess the extent of household responses to school feeding programs. Jacoby (2002) for instance, tests for a 'fly-paper' effect in the Philippines by examining the difference in household calorific intake for children on school and non-school days. There is, however, a gap in the literature on household responses to school inputs, partly due to tricky sampling issues (more on this below) and partly due to the predominance of the

[3] This is not without restrictions; for a more detailed comment see Footnote ??

[4] On this specific topic, see Kremer ??? who argues that the provision of textbooks in Kenya failed to improve test scores due to higher induced class-sizes.

[5] Hoxby (2000) provides another example of using 'surprises' in the context of estimating the effect of peers on cognitive acheivement.

[6] We also follow a close parallel literature on consumption and health. For instance, optimal growth paths derived in our model are similar to those in Foster (1995) who considers the impact of rice prices on child weight in Bangladesh and Dercon and Krishnan (2000) who relate adult sickness to weight in Ethiopia.

production function approach in the literature relating schooling inputs to outcomes. By providing estimates on the size of these responses, we thus suggest areas for future research.

3. Theoretical Framework

3.1. The Conceptual Experiment

Consider a household that receives a single school input, that can either be anticipated or unanticipated. The anticipated input is fully incorporated into the utility maximization problem. For the unanticipated input, households expectations at time $t-1$ (when household decisions are made) are zero, so that they are unable to respond by adjusting their own expenditures. How do these different types of inputs affect cognitive achievement?

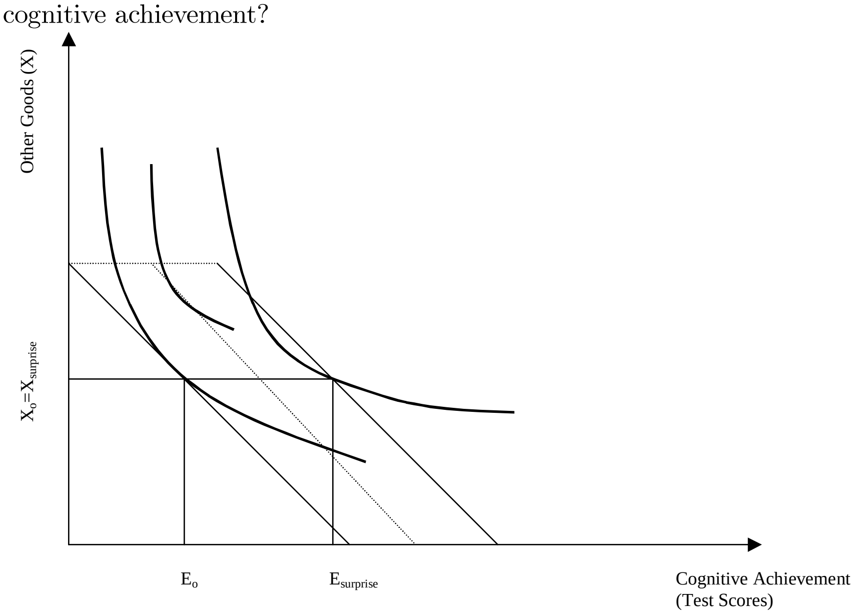

Figure 1 depicts this conceptual exercise in the case where households preferences are defined over cognitive achievement and other goods. For simplicity, we assume that cognitive achievement is related to schooling inputs through a single-input linear production function. In this framework, any additional *unanticipated* input will sustain expenditure on all other goods at X_0. Consequently the change in cognitive achievement will reflect entirely the characteristics of the production function mapping inputs to attainment, moving the household (for instance) from E_0 to $E_{surprise}$. Consider now the impact of inputs of the same magnitude that are fully anticipated by the household at time $t-1$. To the extent that this change is viewed as 'permanent' by the household, the budget constraint will shift out to the outer (non-dashed) linear at time $t-1$ *itself* reflecting the fact that households will optimally incorporate all future information into their decisions at $t-1$. More generally, it could be that the change is (rationally) expected to last a fixed number of years, in which case the budget constraint would shift to an intermediate level (shown by the dashed line) representing the change in permanent income secondary to the anticipated inputs. In either case, there will be no difference in educational inputs provided by the household between $t-1$ and t.This difference forms the basis of our statistical test: Cognitive achievement should respond to unanticipated rather than anticipated inputs.

This special case implicitly assumes that household and school inputs are fully substitutable in the production function for cognitive achievement and therefore anticipated changes have no impact between two subsequent time periods. In the formalization of the model below we introduce two new components. First, we view cognitive achievement as a durable good with household preferences defined over it's stock. Second, we incorporate the production function for cognitive achievement as a constraint in our optimization. Below we see that this affects the program directly through the user-cost of the durable good. In general, anticipated and unanticipated inputs will have differential effects, but the relative size of these effects will depend on the extent of substitutability between household and school provision of the input.

We start with two general assumptions on preferences and the production function for cognitive achievement. An Euler equation is derived that defines conditions governing the growth of test scores—this development relates closely to the discussion on durable goods and inter-temporal household optimization discussed, for instance, in Deaton and Muellbauer (1980), Jacoby and Skoufias (1997) and Foster (1995). Based on this solution we discuss the differential impact of anticipated and unanticipated school inputs on test-score growth. Finally, we consider how credit constraints can affect our results.

3.2. Model

A household (with a single child attending school) derives (instantaneous) utility from the cognitive achievement of the child TS and the consumption of other goods X. The household maximizes an inter-temporal utility function $U(.)$, additive over time and states of the world with discount rate $\beta(<1)$ subject to an inter-temporal budget constraint (IBC) relating assets in the current period to assets in the previous period, current expenditure and current income. Finally, cognitive achievement is determined by a production function relating current achievement (TS_t) to past achievement (TS_{t-1}), household educational inputs (z_t), school inputs (w_t), non time-varying child characteristics (μ) and non time-varying school characteristics (η). We impose the following structure on preferences and the production function for cognitive achievement:

[$A1$] Household utility is additively separable, increasing and concave in cognitive achievement and other goods.

[$A2$]The production function for cognitive achievement is given by $TS_t = F(TS_{t-1}, w_t, z_t, \mu, \eta)$ where $F(.)$ is concave in its arguments.

Under [A1] and [A2] the households problem is

$$\begin{aligned} Max_{(X_t,z_t)}\ U_\tau &= E_\tau \sum_{t=\tau}^{T} \beta^{t-\tau}[u(TS_t)+v(X_t)]\ s.t. && (1)\\ A_{t+1} &= (1+r).(A_t+y_t-P_tX_t-z_t) && (2)\\ TS_t &= F(TS_{t-1}, w_t, z_t, \mu, \eta) && (3)\end{aligned}$$

Here u and v are concave in each of its arguments. The inter-temporal budget constraint (2) links asset levels A_{t+1} at $t+1$ with initial assets A_t, private spending on educational inputs z_t, income y_t and the consumption of other goods, X_t. The price of educational inputs is the numeraire, the price of other consumption goods is

P_t and r is the interest rate. Finally, the production function constraint (3) dictates how inputs are converted to educational outcomes, but we postpone discussion regarding the form of the production function till later.[7]

In this formulation credit markets are perfect so that there are no bounds on A_{t+1}, except that A_{T+1} is equal to zero; all loans need to be paid back and there is no bequest motive. The perfect credit market assumption is relaxed in our discussion on the impact of liquidity constraints below. Moreover, households choose only the levels of X_t and z_t so that school inputs (w_t) are beyond its control. At the time the household has to make its decision, it only knows the underlying stochastic process of w_t and not the actual level. In other words, we assume that school inputs are a source of uncertainty in the model - for simplicity the only source. This assumption is retained throughout the theoretical discussion, but is later relaxed in the empirical test, where we allow for unobserved time-varying characteristics of the household that may influence school inputs.

Maximization of 1 subject to 2 and 3 provides a decision rule related to TS_t, characterizing the demand for cognitive achievement. To arrive at this decision rule, we define a price for cognitive achievement as the 'user-cost' of increasing the stock in one period by one unit, i.e. the relevant (shadow) price in each period for the household. Once such a price is defined, the program is transformed into a standard consumer optimization problem (see for instance, Deaton and Muelbauer, 1980).

To define this price, note that if cognitive achievement could be bought and sold in an open market at the price, v, households would pay v in the first period to buy one unit. In the next period, they could then sell $(1-\delta)$ units (if depreciation is at the rate $1-\delta$) in the second-hand market and receive the current value $\frac{(1-\delta)v}{1+r}$. Thus, the cost of holding one unit of test-scores for one period is $v - \frac{(1-\delta)v}{1+r}$ and this defines the user cost. In the context of a production function however, households 'buy' test-scores in period t is by increasing z_t. Since we are interested in the cost of boosting achievement in one period only we assume that in the next period they can reduce z_{t+1} to ensure that the overall stock of test-scores at $t+1$ remains unchanged. The user cost, evaluated at period t, is then (see Appendix 1 for the derivation):

$$\pi_t = \frac{1}{F_{z_t}(.)} - \frac{F_{TS_t}(.)}{(1+r)F_{z_{t+1}}(.)} \tag{4}$$

Similar to the expression $(v - \frac{(1-\delta)v}{1+r})$ derived above, the first term measures the cost of taking resources at t and transforming it into one extra unit of cognitive achievement. When implemented through a production function, the price is no longer constant— if the production function is concave, the higher the initial levels of cognitive achievement, the greater the cost of 'buying' an extra unit as reflected in the marginal value,

[7] There are two observations regarding the form of the utility function. First, an alternative assumption, that the benefits from the child's cognitive achievement are only felt in the future, would not change the model fundamentally. If these benefits are only related to the flow of earnings in the future from the child's cognitive achievement, then the education decision becomes similar to a pure investment decision. As long as the benefits from education are concave in its arguments, the results would be similar. Note that this of course, does not imply that the steady state value of human capital will be the same in either case, but only that along the growth path first order conditions remain unchanged (see Banerjee, 2003 for a detailed discussion of steady states under different assumptions regarding the form of the utility function). Second, the utility function uses a stock as one of its arguments: We assume that households care about the level of educational achievement. The results below are unaffected if one assumes that households care about the (instantaneous) flow from educational outcomes, provided that this flow is linear in the stock.

$F_{z_t}(.)$. Of the additional unit bought in period t, the amount left to sell in period $t+1$ is $F_{TS_t}(.)$ and the second term thus measures the present value of how much of this one unit will be left in the next period expressed in monetary terms. Once this is defined, the standard first-order Euler condition related to the optimal path education of educational outcomes between period $t-1$ and t can be derived as:

$$E_{t-1}\left(\beta\frac{\pi_{t-1}}{\pi_t}\frac{\frac{\partial U}{\partial TS_t}}{\frac{\partial U}{\partial TS_{t-1}}}\right)=1 \tag{5}$$

Intuitively this expression (ignoring uncertainty for the moment) suggests that if the user-cost of test-scores increases in one period t relative to $t-1$, along the optimal path this would increase the marginal utility at t, so that TS_t will be lower. This is a standard Euler equation stating that along the optimal path, cognitive achievement will be smooth, so that the marginal utilities of educational outcomes will be equal in expectations, appropriately discounted and priced. Finally, the concavity of the production function will limit the willingness of households to boost education fast since the cost is increasing in household inputs. Thus, under reasonable restrictions, the optimal path will be characterized by a gradual increase in educational achievement over time (for an explicit derivation of the Euler equation with durables, see e.g. Deaton and Muellbauer (1980), Foster (1995)).

To proceed with the empirical specification impose the following conditions on preferences and the production function.

$[\widehat{A1}]$ Household utility is additively separable and of the CRRA form.

$[\widehat{A2}]$ $TS_t=(1-\delta)TS_{t-1}+F(w_t,z_t,\mu,\eta)$ where the Hessian of $F(.)$ is negative semi-definite.

Under $[\widehat{A1}]$ marginal utility is defined as $TS_t^{-\rho}$, with ρ the coefficient of relative risk aversion. Then, 5 can be rewritten as:

$$\left(\frac{TS_t}{TS_{t-1}}\right)^{-\rho}\frac{\beta\pi_{t-1}}{\pi_t}=1+e_t \tag{6}$$

where e_t is an expectation error, uncorrelated with information at $t-1$. Taking logs and expressed for child i, we obtain:

$$\ln\left(\frac{TS_{it}}{TS_{it-1}}\right)=\frac{1}{\rho}\ln\beta-\frac{1}{\rho}\ln(\frac{\pi_{it}}{\pi_{it-1}})+\frac{1}{\rho}\ln(1+e_{it}) \tag{7}$$

or, the growth path is determined by the path of user costs, and a term capturing expectational surprises.

3.2.1. Anticipated and Unanticipated Inputs

A key issue is how increases in school level inputs w_t impact on the optimal path of cognitive achievement. Since school resources are not known with certainty until after households make decisions regarding their own inputs, this impact will depends on whether such increases are anticipated. Thus, let $w_t^a(w_t^u)$ be inputs at time t that were anticipated (unanticipated) at $t-1$. First, consider increases that are anticipated. In this case, the impact on the path of outcomes will depend on its impact on the user cost of educational achievement at t, since there is no *direct* impact on the budget constraint (all information included anticipated inputs will have been incorporated into the budget constraint at time $t-1$). In particular, using the implicit function theorem with 4 and assuming [$\widehat{A2}$], we have

$$\frac{d\pi_t}{dw_t^a} = -\frac{F_{z_t w_t}}{F_{z_t}^2} \lesseqgtr 0 \; if \; F_{z_t w_t} \gtreqless 0 \tag{8}$$

This implies that if household and school inputs are technical substitutes ($F_{z_t w_t} < 0$), anticipated increases in school inputs at t will increase the relative user cost of boosting at t, resulting in lower growth of cognitive achievement, ceteris paribus, between t and $t-1$, consistent with the optimal path (5). Alternatively, if school and households inputs are technical complements, increases in school inputs at t will increase the marginal productivity of household inputs at t, and through the decline in user costs lead to higher growth in cognitive achievement along the optimal path between t and $t-1$.

To clarify the dynamics between $t-1$ and t further, note that there are two effects we need to distinguish. The first is due to the change in relative user costs while the second is governed by the households desire to smooth consumption (5). The second effect will always provide incentives to spend more at $t-1$ to take advantage of the additional government spending at t. If household and school inputs are substitutes, households will optimally recognize that relative user costs at t will be higher than at $t-1$- the 'implicit price' of buying test-scores will increase in the future. Consequently, to retain the optimal growth path of 5 households will choose to increase their own spending at $t-1$. Thus, the growth in cognitive achievement will be lower relative to the case where no household responses are possible.

In the case of technical complements, the behavioral response is exactly the opposite- since relative user costs will be lower at t, households will optimally delay spending. However, in this case the user-cost and the smoothing effects move in opposite directions so that the overall growth could still be higher relative to the case where $w_t = 0$. Comparing the two cases of complements and substitutes, household spending is thus counter-cyclical relative to government spending when household and school inputs are substitutes. When they are complements, the smoothing and the user-cost effects move in opposite directions, although in the special CRRA case that we consider here ($\widehat{A1}$) the user-cost effect is higher than the smoothing effect so that the pro-cyclicality of household inputs is maintained.

For *unanticipated* increases in school inputs,. since households are unable to respond, they are pushed off their optimal path and the increase in educational achievement in period t is given by $F_{w_t} dw_t$. What is the size of this effect compared to a similar increase in anticipated inputs? When inputs are anticipated, using 7 the change in the optimal growth path is given by

$$\begin{aligned} \frac{\partial(\Delta_{t-1}^t \ln TS)}{\partial w_t^{anticipated}} &= -\frac{1}{\rho}\left(\frac{\partial \ln \pi_t}{\partial w_t}\right) \\ &= \frac{1}{\rho}\frac{1}{\pi_t}\frac{F_{zw}}{F_z^2} \lesseqgtr 0 \; if \; F_{zw} \lesseqgtr 0 \end{aligned} \tag{9}$$

For unanticipated increases the change in the growth path is give by $\ln(TS_t + w_t^{unant} F_w)$ which is strictly positive. Thus, the effect of an unanticipated change is higher than that of an anticipated change in the case of substitutes, and relative sizes cannot be ranked when they are complements without further restrictions on the form of the utility function. These results are summarized in the table below for the case of the CRRA.

Type	**Inputs**	**Cross-Derivative**	**Spending at t-1**	**Spending at t**	**Effect on Growth**	**Relative Ranking**
A	Anticipated	Substitutes	Increases	Decreases	Lower	A<{B,C,D}
B	Anticipated	Complements	Decreases	Increases	Higher	B>A
C	Unanticipated	Substitutes	Unchanged	Unchanged	Higher	C>A ;C=D
D	Unanticipated	Complements	Unchanged	Unchanged	Higher	D>A; C=D

We stress here that increases in outcomes due to unanticipated inputs (in the case of substitutes) are *sub-optimal*; household education spending will be higher than that justified by the decline in the user cost of boosting educational achievement. Consequently, in the next period, the household will implement a correction using the 'correct' user cost to restore themselves to the optimal path. Cognitive achievement being a durable however implies that the increase in outcomes will not be entirely undone, since incentives will exist to sustain the stock at a higher level than before. While we return to this in our discussion below, it is important to realize therefore that the distinction between anticipated and unanticipated inputs has greater relevance for identification purposes than it has for policy—although unanticipated inputs lead to larger increases in cognitive achievement in the current period, this will be smoothed out subsequently.

3.2.2. *Credit Constraints*

A straightforward way to incorporate credit constraints is to assume that any point in time assets have to be nonnegative, i.e. credit is impossible unless fully collateralized ($A_t \geq 0$). The definition of the user cost remains unaffected so that the main impact is that credit constraints may limit the ability to equate appropriately discounted and priced marginal utilities. More specifically, let λ_t be the multiplier linked to the non-negativity constraint of carrying over assets between $t-1$ and t, then the optimal path, linking $t-1$ and t, can be defined as:

$$\frac{\partial U}{\partial TS_{t-1}} = E_{t-1}\left(\beta \frac{\pi_{t-1}}{\pi_t} \frac{\partial U}{\partial TS_t}\right) + \lambda_t \tag{10}$$

in which a binding credit constraint ($\lambda_t > 0$) would result in higher marginal utility at $t-1$ (lower educational achievement), than what would have been implied without credit constraints ($\lambda_t = 0$)—see also 5. What is the impact of unanticipated and anticipated changes in w_t in this case? First, the impact of unanticipated changes in unaffected: since no behavioral response is possible, the effect still works through the production function of cognitive achievement. A difference in this case is that with binding credit constraints the impact may be to bring the household closer to its unconstrained optimal path by alleviating the credit constraint (since fewer private resources would be needed in the 'optimal' case, relaxing the budget constraint). The impact of anticipated changes now works via two effects: First the effect through the change in user costs has exactly the same impact as before. Second, there is a further effect through the budget constraint: Since fewer household resources are required than before to achieve the same level of educational achievement, this would alleviate the budget constraint (i.e. reduce the shadow cost of the constraint, λ_t).

The size of the overall effect on cognitive achievement would depend on the income effects related to the alleviation of the budget constraint. If the reduction in credit constraints leaves spending on private inputs unaffected (so that no more other goods are consumed despite the relaxation of the budget constraint), then the impact would be indistinguishable from unanticipated changes. However, in the more plausible case that other goods are normal, in general, the effect of the anticipated relative to unanticipated changes will retain the same ordering (i.e. anticipated changes have a lower (higher) effect if school and household inputs are substitutes (complements) caveated, as usual, by our discussion on the relative ranking of growth effects in the case of complements.

3.2.3. Summary

The key results of this section can thus be summarized as follows:

1. When household and school inputs are substitutes, an increase in anticipated inputs at t will lead to an increase in household inputs at $t-1$, a decrease at t and subsequently a lower rate of growth of cognitive achievement.
2. When they are complements, the opposing signs of the 'user-cost' and the 'smoothing' effects imply that overall effect depends on household preferences.
3. Unanticipated inputs on the other hand, have no impact on household inputs at t and $t-1$ and thus always lead to a higher growth in cognitive achievement between t and $t-1$.
4. Finally, these results remain unchanged (though attenuated) by the imposition of credit constraints under the mild assumption that other goods consumed by the household are normal in the sense that their income-elasticity is positive.

In testing the predictions from this model using cash grants as the relevant input, we recognize that schools should optimally allocate these grants across different inputs. One way to interpret these results is that schools are constrained in what they can do and are hence unable to spend cash grants optimally. These constraints could arise either due to thin markets (for instance, in the case of teachers) or lack of scale economies (for instance, to improve infrastructure).[8] The estimated equations are thus a 'reduced-form' in the sense that they represent the effect of grants taking into account constrained maximization at the school level.

4. Empirical Specification

Our empirical strategy is based on two related tests. We first test whether household educational expenditures and school cash grants are substitutes. The theory implies that if this is the case, contemporary household funding z_t should decline with an increase in w_t. This is a cross-sectional test- as long as the classical assumptions of ordinary least squares are maintained, we should find that households matched to

[8] As examplex, in the case of teachers, while most head-teachers complained of shortages, only in 2 cases were teachers hired by the school. Both turned out to be significantly worse (in terms of education and training) than government teachers, leading to considerable dissatisfaction among the community.

schools with higher cash grants spend less on their children's education. Once we have established that household and school inputs are technical substitutes, we proceed to examine the hypothesis that (under the assumption of technical substitution) the impact of anticipated grants is smaller than that of an equivalent unanticipated amount. We detail each of these in turn.

4.1. Testing Substitutability between Household and Cash Grants to Schools

We estimate a generic demand model in which household spending on school-related inputs is regressed on wealth (proxying permanent income), school attributes and school grants

$$\ln(z_i) = \alpha + \beta_1 A_i + \beta_2 \ln w^a_{j(match\ i)} + \beta_3 \ln w^u_{j(match\ i)} + D_i P_i + \varepsilon_i + \varepsilon_j \tag{11}$$

where z_i is the spending by the household on child i, A_i is the assets owned by household of which i is a member (as a proxy variable for the permanent income of the household), w^a_j is anticipated grants in school j that matches to child i, w^u_j is unanticipated grants received by school j and P_i are province level controls. The error term is decomposed into two components where ε_i is a child specific and ε_j, is a school specific error term. We test $\beta_2 < \beta_3 = 0$, i.e., households respond negatively to expected grants at the school level by cutting back their own funding, but are unable to respond to cash grants that are unanticipated.

The estimation problem arises from an omitted variable bias if households exercise school choice (either through the choice of a particular school or the decision to keep the child at home). In particular, we will see that in the Zambian case, $w^a_{it} = \frac{\text{Constant}}{\text{enrollment}}$, so that an omitted variable that increases the probability of sending a child to a specific school as well as the *unconditional* (on school choice) spending on educational materials at the household level will lead to a correlation between the error term and the school level per-pupil grants. Further, since the choice of school will depend on the vector of attributes for *all* schools that the household could potentially choose from, a survey would have to include all *potential* schools in the catchment area of households. Moreover, to control for this omitted variable problem, an explicit selection equation for school choice would have to be estimated at the first stage.

As an alternative to this strategy we base our identification on careful sampling taking advantage of the very high historical enrolment rates in Zambia, even in remote rural areas. Thus, we restrict attention only to those villages where the distance costs of travelling to a school other than the one surveyed are very high (more on this below) so that any potential benefits of choosing an alternative school are unlikely to outweigh the cost of transportation. Under this assumption the problem is simplified to a two-dimensional choice between schooling and no-schooling. Using this sample of households we then show that enrollment in schools is exogenous to child level educational expenditures and presumably arises from differences in the cohort size of the catchment area. Our strategy of using variation in cohort size to examine the impact of schooling inputs has been previously exploited by Case and Deaton (1999) and Urquiola (2003). We provide stronger evidence for the validity of this strategy by explicitly testing the crucial assumption that enrollment is exogenous to our dependent variable.

From this basic hypothesis, we can further exploit the data to test whether the grants are 'truly' anticipated in the sense that households make their own decisions taking anticipated cash grants into account even before such grants are actually received. If there are schools where, at the time of the survey, $w^a \neq 0$ but $w^a_{received} = 0$, a test can be based on the difference in the estimate of $\widehat{\beta_2}$ depending on whether we use w^a or $w^a_{received}$ as the explanatory variable.[9] Specifically if household decisions are based on anticipated rather than received grants, $\widehat{\beta_2}(w^a) > \widehat{\beta_2}(w^a_{received})$ since if $w^a > 0$ household spending will be less than what would be predicted by using $w^a_{received}(= 0)$ as the explanatory variable. Thus, we can estimate a second equation

$$\ln(z_i) = \alpha + \beta_1 A_i + \widetilde{\beta_2} \ln w^a_{j(received)} + \beta_3 \ln w^u_j + D_i P_i + \varepsilon_i + \varepsilon_j \tag{12}$$

and test $\widetilde{\beta_2} < \beta_2$. A rejection of the null would lead to greater confidence in the technical-substitution results since it would not only imply that cash grants to schools crowd-out household funding ($\beta_2 < 0$), but further that households *anticipate* such grants and make their expenditure decisions before the grants are actually disbursed.

Our final specification then breaks up the household response by asset terciles- in particular, if credit constraints are important, we would expect poorer households to respond *less* to expected grants than the rich. Thus we estimate

$$\ln(z_i) = \alpha + \beta_1 A_i + \widetilde{\beta_2} \ln w^a_j + \beta_3 \ln w^u_j + \beta_4 (A_i * \ln w^a_j) + D_i P_i + \varepsilon_i + \varepsilon_j \tag{13}$$

where the additional term are interactions between the asset index and the expected grants (operationalized in the estimation as interactions between asset terciles and expected grants). Under credit constraints, we would now expect $\beta_4 < 0$ as well, so the elasticity of substitution increases with the wealth of the household.

4.1.1. Robustness to Potential Omitted Variables

The key question of whether the relationship estimated above can be given a causal interpretation depends on the source of variation for the total enrollment in the sampled school: If such variation is due to cohort-size, a causal interpretation is likely (maintaining the exogeneity of cohort size). Alternatively, if this variation arises due to differential probabilities of enrollments among household in the catchments of the surveyed school, our coefficient could be biased upwards- it is plausible that households who choose higher levels of schooling also incur larger expenditures on their children's schooling.

In order to tackle this, we begin by treating the household expenditure function as a censored regression and estimate it as a Tobit model in the standard fashion, including all children of school going-age in the sample. However, this procedure, while it takes care of the possibility that enrollment and spending are driven by the same mechanism, does not tackle the kinds of issues raised by the possibility that enrollment might be endogenous to spending. To examine whether this might be so, we use the method proposed by Smith and Blundell(1986). We first estimate the determinants of anticipated funding (recall this is K/enrollment) in school j, using the average distance to the school for households in the village, the distance of the nearest catchment area boundary (interacted with a dummy for rural schools), the number of primary school age going children in the sample and the average wealth index in the school (as a proxy for catchment

[9] Satisfying the requirement that $w^a = 0$ at the time of the survey is uncorrelated to the error term of the regression. Our identification is based on the fact that these were schools that were surveyed earlier in the month combined with delays in disbursement.

area wealth). We obtain the residuals from this equation and re-estimate the Tobit model of household expenditures, including the estimated residuals. The t-statistic on the estimated coefficient of the residuals provides a test of weak exogeneity for anticipated funding. The identifying assumption then, is that the placement of public schools as well as cohort size is exogenous to household educational expenditures.

4.2. Test-Score Hypothesis

Once we have established that household and school grants are technical substitutes, we can restrict ourselves to testing whether the impact of unanticipated grants on gains in cognitive achievement is higher than that of anticipated grants. A major concern in the parallel literature on production functions when testing for the relationship between school inputs and cognitive achievement has been the presence of unobserved child and school level heterogeneity. In the context of the Euler framework a fairly flexible form of heterogeneity can be incorporated into the production function without biasing the estimated coefficients. To see this return to equation 7 and the production function given by $[\widehat{A2}]$ $(TS_{it} = (1-\delta)TS_{it-1} + F(w_t, z_t, \mu_i, \eta_i))$. We showed previously that $\pi_t = \frac{1}{F_{z_t}(.)} - \frac{(1-\delta)}{(1+r)F_{z_{t+1}}(.)}$. Now, as long as $F_{z\mu_i} = 0$ *or* $F_{z_t} = \mu_i(g(w_t, z_t))$, so that $\ln \pi_{it} = \ln \pi_t \eta_i$, the unobserved heterogeneity embodied in μ_i is eliminated from the estimating equation. This formulation is then sufficient to ensure that the path of user costs is defined only in contemporaneous variables and is unaffected by fixed heterogeneity and past school achievement.[10] Assuming identical risk preferences, an empirical specification consistent with [13] can then be written as:

$$\ln\left(\frac{TS_{it}}{TS_{it-1}}\right) = \alpha_o + \alpha_1 \ln w_{it}^a + \alpha_2 \ln w_{it}^u + \alpha_3 \Delta X_t + \epsilon_{it} \tag{14}$$

in which w_{it}^a and w_{it}^u are anticipated and unanticipated input changes, in this paper proxied by flows of funds, while ΔX_t reflects all other sources of changes in the user cost between t and $t-1$. The core prediction is that the marginal effect of anticipated is lower than unanticipated funds when household and school inputs are substitutes. This prediction is unaffected by the presence of credit constraints, even though α_1 is likely to be larger in that case.

[10] How restrictive is this particular formulation of the production function? Note first that we can either write the production function as

$TS_{it} = (1-\delta)TS_{it-1} + \beta_1 w_t + \beta_2 z_t + \beta_3 \mu_i + \beta_4 \eta_j$

or as

$TS_{it} = (1-\delta)TS_{it-1} + f(w_t, z_t, \mu_i \eta_j)$ if we make sure that $F_{z_t} = \mu_i(g(w_t, z_t))$. We can compare this to three popular specifications used in the literature on the estimation of production functions for cognitive achievement, discussed in Todd and Wolpin (2003). The first, the contemporaneous specification has no role for either past levels of cognitive achievement or for the possibility that household inputs z_t will be correlated to unobserved school and child level variables. The production function that we use here allows for both of these possibilities. The second specification, which has been widely used in the recent past is the value-added specification

$T_{it} = \beta_1 T_{it-1} + \beta_2 w_t + \beta_3 z_t$

Note that in this case, the unobserved heterogeneity is assumed to enter only at time 0 so that $T_{i0} = \mu_i$- the child's mental 'endowment' leads to a fixed increase in test scores, instead of an incremental increase in every period. A more general form is given by the cumulative specification where

$T_{it} = \beta_1 T_{it-1} + \beta_2 w_t + \beta_3 z_t + \beta_4 \mu_i$

so that child endowment can affect cognitive achievement in every period. These three widely used specifications have increasing data requirements. In particular, the cumulative specification would require at least 3 periods of data to arrive at consistent estimates- a tall order in a low income country context. Our specification of the production function not only allows for the cumulative specificaiton, but also allows unobserved heterogeneity to enter in a multiplicative form so that the marginal value of household inputs can depend on unobserved child and school endowments. Further, the production function also allows for the possibility that past inputs may affect current cognitive achievement through ways other than lagged achievement- specifically, as long as we maintain the additive separability of F_{z_t} and past inputs, our user costs will remain unaffected. The immense flexibility in the form of the production function for cognitive achievement that the Euler framework allows for is then a major advantage over attempts to directly estimate such relationships..

There first econometric concern relates to the identification of 'anticipated' and 'unanticipated' cash grants. Below, we will see that 'anticipated' grants are well identified—such grants are based on a legislated rule and a detailed tracking (implemented in the survey itself) confirms that schools receive exactly the amount stipulated. Our interpretation of unanticipated grants may be more problematic since time-series data on cash grants which could be used to calculate deviations from the mean are not available. Specifically, we assume that grants other than the anticipated amount, which are determined at the discretion of the District Education Office are unanticipated. This would probably overstate the unanticipated component-schools could have been informed previously of such grants, or there could be differential expectations for different schools. To see how this affects our estimates, ideally we would like to estimate

$$\ln\left(\frac{TS_{it}}{TS_{it-1}}\right) = \alpha_o + \alpha_1 \ln w_{it}^a + \alpha_2 \ln w_{it}^u + \varepsilon_{it} \tag{15}$$

but (assuming w.o.l.o.g for the moment that observed anticipated grants are zero) our estimating equation is actually

$$\ln\left(\frac{TS_{it}}{TS_{it-1}}\right) = \alpha_o + \widehat{\alpha_2} \ln(w_{it}^a + w_{it}^u) + \varepsilon_{it} \tag{16}$$

Under the assumption that $\alpha_2 > \alpha_1$, $\widehat{\alpha_2} < \alpha_2$—our estimates of the impact of unanticipated grants constitute a lower bound, where the extent of attenuation depends on the degree to which our construction of unanticipated grants may actually contain components that were anticipated. To see this formally we can write the mis-specification as

$$\begin{aligned} True\ Model &: \ y = b_1 X_1 + b_2 X_2 + \varepsilon \\ EstimatedModel &: \ y = \beta(X_1 + X_2) + \varepsilon \end{aligned}$$

where we are interested in the relative size of $E(\widehat{\beta})$ compared to b_1 and b_2 under the assumption that $b_1 > b_2$. Then,

$$\begin{aligned} \widehat{\beta} &= \frac{\sum(X_1+X_2)y}{\sum(X_1+X_2)^2} \ Let\ X_1, X_2 \text{ have 0 mean w.o.l.o.g} \\ &= \frac{\sum(X_1+X_2)(b_1X_1+b_2X_2+\varepsilon)}{\sum(X_1+X_2)^2} \\ &\Rightarrow \ E(\widehat{\beta}) = \frac{b_1 Var(X_1) + (b_1+b_2)cov(X_1,X_2) + b_2 var(X_2)}{var(X_1)+var(X_2)+2cov(X_1,X_2)} \end{aligned}$$

and a sufficient condition for $b_1 > E(\widehat{\beta}) > b_2$ is that $cov(X_1, X_2) \geq 0$. This covariance is necessarily zero if X_1 and X_2 are unanticipated and anticipated components respectively.

The second concern is potential bias due to non-zero correlation between the error term and unanticipated cash grants arising from dynamic heterogeneity (time varying school or district specific effects).[11] This may be the case for instance if there is a change in a school level variable that leads, on the one hand, to greater unanticipated grants and on the other to higher learning gains (the introduction of a highly motivated teacher

[11] Note though that omitted variable bias due to heterogeneity of fixed characteristics, which normally plague cross-sectional estimates, are accounted for through the Euler framework.

who both searches for such funds and teaches exceptionally well is an example). With such omitted variables, our estimate of the impact of unanticipated grants would be biased upwards.

We correct for this problem through an instrumental variables strategy. From the Euler framework, any variables at $t-1$ are available for use as instruments since such information will have been incorporated into the decision making process. In addition we also use variables that were unknown at $t-1$ but had an impact on unanticipated grants at t. In particular, based on the detailed tracking of funds, we use per-pupil grants that the *district* office received from external (non-governmental) sources that was allocated without any consultation with the offices.

For our instrumentation strategy to be valid, we require that the instrumental variables are positively correlated with the amount of unanticipated funds received, but are not correlated with the *gain* in cognitive achievement over the year. First, the amount of such grants boosts the overall funds available at the district level and hence the unanticipated funds passed on to schools. Simultaneously, since such funds arrive from external sources, it is unlikely that districts were able to actively influence the amount of cash grants that they would receive. This addresses our main concern that there might have been changes at the district level such that districts that received more grants were also more likely to 'place' this in a targeted manner. Consequently, our strategy of using district level aggregates combined with interactions of lagged stock variables at the school level generates predicted cash grants that are uncorrelated to the error term in the regression above.

Finally, our results are reported at the school level. Since this is a straightforward linear aggregation, there is no reason to expect results to change; simultaneously, we are better able to handle the clustering of errors at the level of the district. To ensure that the sample remains the same, we compute school level scores only for those students who were present in both years for the test.

4.3. Other Econometric Concerns

There are two other specific concerns that arise due to the specific nature of our data. We briefly detail our strategy in dealing with each of these below.

4.3.1. Treatment of Zero Cash Grants

The first is that a large number of schools receive zero unanticipated grants in the sample. Since our estimation equation is based on the log transformation, we need to modify this variable in order for the log to be defined. Moreover we need to address this problem both when cash grants are explanatory (in estimations of the impact on cognitive achievement) as well as dependant variables (in the first stage of our instrumentation strategy). We address each of these cases in turn.

We use a modification of the method developed by Johnson and Rausser (1971) to derive the optimal constant to be added on to zero values. The basic intuition behind this approach is that the constant should be chosen so that the estimated relationship between cash grants and cognitive achievement is identical for both schools with zero and non-zero cash grant values (dealing with potential selection issues through the

IV strategy above). In particular, we can treat the sample of schools as two separate samples consisting of m observations of zero grant values and $n - m$ observations of positive grant values. The regressions can then be represented as

$$\begin{aligned}
\ln(\Delta TS_i) &= \delta_1 \ln(X_i + b) + \varepsilon_i \\
i &= 1, 2, ..., m \\
\ln(\Delta TS_i) &= \delta_2 \ln(X_i) + \varepsilon_i \\
i &= m + 1, ..., n
\end{aligned}$$

and b is estimated under the restriction that $\delta_1 = \delta_2$. In addition to presenting estimated coefficients based on the estimated $\widehat{b}$, we also present robustness tests based on estimated coefficients when the unanticipated grants are treated as dummy variables (i.e. making a distinction only between those who received and those who did not). We show that our results are robust to these alternative specifications.

For the case where cash grants are the dependant variable, we estimate a hurdle model (Wooldridge, 2001) where the probability of receiving such grants is estimated separately from the amount received *conditional* on receipt. Under the assumption that the dependant variable, y, is distributed log-normally conditional on $y > 0$, the maximum likelihood estimate of the unconditional $E(y)$ can be shown to be

$$E(y|x) = \exp(x\beta + \frac{\sigma^2}{2})\Phi(x\gamma)$$

where $x\beta$ is the predicted value from the OLS regression of $\ln y = \beta x + \varepsilon$ (restricted to $y > 0$) and $\Phi(x\gamma)$ is the predicted probabilities estimated from a probit. In this instance, a hurdle model is preferable to the Tobit since the latter requires the probability of receipt as well as the amount received (conditional on receipt) to be governed by the same process. Contrary to this, the hurdle allows for these processes to differ, so that in the predicting equations the process of determining which schools receive positive grants is separate from the determination of how much the schools receive. Note that since we are interested in the $E(\ln y|x)$ and not $\ln(E(y|x))$ the above simplifies to

$$E(\ln y|x) = (x\beta)\Phi(x\gamma)$$

Finally, the prediction for this hurdle model is then used in the second stage of our instrumentation strategy, with an appropriate standard error correction for the use of generated regressors (Murphy and Topel, 1985).

4.3.2. Measurement Error: Test Scores

Our second concern relates to the treatment of test-scores. We model test scores as arising from the distribution of an underlying latent variable (θ) following the literature on Item Response Theory (Birnbaum, 1967). This method has several advantages in that the properties of the estimated latent variables are easy to interpret and the importance of the *characteristics* of the test are made explicit in the estimation. Further, the maximum likelihood procedure used to estimate the latent variable generates weights that are 'locallyoptimal' in the sense that they minimize the error of classification. The latent variable, θ, is estimated through a maximum likelihood procedure using a structural assumption regarding the mapping between θ and the probability of a correct response. The standard error of this estimate can then be computed as $\frac{1}{\sqrt{\sum_j I_j(\theta)}}$ where $I_j(\theta)$ is Fisher's information for a particular question, j (the procedure is detailed in Appendix 1).

How does this error of measurement effect our estimates? To the extent that test-scores are the dependant variable in our regressions, this increases the standard error of the regression, but our estimates remain consistent and unbiased. A correction is required however, since the error structure is now characterized by a variance-covariance matrix that violates homoskedasticity of the disturbance terms. Since the standard error of our estimate is a function of θ, the variance-covariance matrix consists of terms like $\sigma_{\varepsilon}^2 + \sigma_{u_i}^2$, where σ_{ε} is the regression error, and σ_{u_i} the measurement error for individual i. We account for this by adjusting standard errors for an arbitrary error-structure due to clustering.[12]

A more serious problem arises if cash grants are targeted towards poorly performing schools in the base year. Measurement error in test scores implies that gain-scores are higher for schools that performed poorly in the base year (Kane and Staiger, 2001 and 2002). Thus improvements in cognitive achievement could arise due to mean reversion rather than a causal relationship with such grants. Although measurement error due to differences in cohorts (a major source of variation in Kane and Staiger, 2002) are eliminated by retesting the same children we do find that in the case of Mathematics the gradient between gains and initial scores is significantly negative. To address this issue, we first show that there is no evidence of targeting towards poorly performing schools and we then instrument for potential endogenous placement using the strategy detailed above.

5. Data

5.1. The Country Context: Education in Zambia

Zambia is a landlocked country with a population of 10 million, almost entirely dependant on copper for export revenues. With a decline in copper prices, there has been a commensurate decrease in income and government resources. As a result, average real per capita government education expenditure in 1996-98 was only about 73 percent of its 1990-92 level, declining further to an average of about 60 percent of this level by 1999-2000 (World Bank data based on Government of Zambia Financial Statements).

This economy wide decline has also had an impact on educational attainment. Although Zambia outperforms other African countries with similar per capita income levels, net primary school enrolment at about 72 percent is at a historically low level, having seen some decline during the previous decade (Figures 2a and 2b).[13] Both the government and households have responded to this worsening of the education profile.

---

[12] Greater efficiency can be obtained by using the standard error of the estimates to implement a (modified) variance weighted least squares estimation. However, simulations in Das and DeLaat (2003) show that the efficiency gains from doing so (compared to the robust sandwich estimator) are small- the major efficiency gains arise from the use of estimated latent variable rather than the test score itself.

[13] These levels are similar to Kenya, higher than Congo or Mozambique, but below those typically attained in other Southern African countries (see for example, World Bank and UNESCO).

The government for instance initiated a Basic Education Sub-sector Investment Program in 2000, which along with administrative changes in the delivery of educational services and restructuring of the teachers payroll also led to some direct financing of schools through cash disbursements. While household responses are clearly harder to interpret, we will see below that parents tend to be active in their children's education with high contributions, both in terms of expenditures as well as time. It is precisely this involvement that will be exploited in our tests below.

5.2. Sampling

Data for the study comes from 182 schools in four provinces of the country and was collected by the authors in 2002[14]. The choice of schools was based on a probability-proportional-to-size sampling scheme, where each of 35 districts in the four provinces was surveyed and schools were randomly chosen within districts with probability weights determined by enrollment in the school. Thus, every *enrolled* child in the district had an equal probability of being in a school that participated in the survey.[15]

In every school, 20 students were randomly chosen in Grade V in 2001 and an achievement test was administered in Mathematics, English and the vernacular.[16] The same test was administered again in 2002 to the same students leading to the construction of a two year panel of test scores. In addition to the tests and a school questionnaire, questionnaires were administered to the head-teacher and all teachers who were teaching or had taught the tested children. These children were also asked to complete a pupil questionnaire in every year with information on basic assets, demographic information and educational flows within the household. Further, as part of the expenditure tracking exercise, district and provincial educational offices associated with the surveyed schools were administered questionnaires detailing financial activity over the year (receipts and disbursements of cash and materials).

In addition to the school survey, household surveys were administered to 540 households in 35 villages. The choice of villages was designed to eliminate complications arising from school choice (see Section II.2): Based on a geographical mapping of all schools, those that satisfied a 'remoteness' criteria (defined as the closest edge of the relevant Thiessen polygon lying at least 2.5 kilometers from the school) were chosen as starting points for villages in the household survey (so that the school was at least 5 kms from the nearest other school). From these schools, the closest (or second closest depending on a random number) village was chosen and 15 households were randomly chosen from households with at least one child of school going age.

Two different samples are thus used for the empirical section of this paper. The first sample (the Household Sample) is based on a subset of 35 remote schools, with data on matched school and household inputs for 540 households. Since 15 households were selected from every village, we have data on cognitive achievement for only 200 students in this sample and hence can use this data only to test the resource-substitution hypothesis. The larger sample of 182 schools provides data on changes in cognitive achievement

[14] Lusaka, Northern, Copperbelt and Eastern provinces were surveyed. These four provinces account for 58% of the total population in Zambia.

[15] Thus, although Zambia has high enrolment rates of above 90%, this survey still may not be appropriate for examining educational outcomes such as enrolment.

[16] In cases with less than 20 students, all children were tested.

for 2,600 students with matched data on school, teacher and head-teacher attributes but not on household expenditures- this sample is used to test the Test-Score hypothesis. Finally, data from provincial and district levels are used to provide instrumental variables for the Test-Score Hypothesis. The table below clarifies the use of our data

Sample	**Questionnaires**	**Learning Achievement**	**Used For**
Household	Household Questionnaire	X	Substitution Hypothesis
School	School Funding, School Attributes	3,600 students in 2001	
	Head Teacher and Teacher Attributes	2,700 re-tested in 2002	Test Score Hypothesis
Tracking	District/Province Level Funding	X	Instrumental Variables

5.3. Description: Schools

Reflecting the overall decline in the education sector, schools in our sample are under some stress (Table 1a): There are over 100 children for every functional classroom, student-teacher ratios are above the Zambian guideline of 40 and there are a large number of repeaters. Moreover, for almost every variable rural areas tend to do worse than their urban counterparts and this difference is further magnified in the case of the 'remote' school sample. The difference is particularly marked for asset holdings where the mean value of the asset index is one standard deviation lower in rural and 1.2 standard deviations lower in remote villages compared to urban areas.

Turning to educational inputs, there are three distinct types of inputs that schools may receive—teaching inputs through new teachers or increases in staff remuneration, in-kind receipts in the form of textbooks or chalk, and cash receipts. The effect of teacher inputs is studied in some detail in Das, Dercon, Habyarimana and Krishnan (2003). Further, during the year of the survey, schools received very little inputs in-kind—on average less than 0.05 textbooks, 0.012 desks, 0.001 chairs and 0.01 boxes of chalk were received per student.[17] Consequently, the impact of the third type of inputs (cash-receipts excluding teacher's salaries) on cognitive achievement forms the basis for this study.

Contrary to the poor record of in-kind receipts, most schools had received some cash and this is explored further in the 'cash-grants' rows of Table 1a. There were two kinds of cash receipts that schools could receive. **Rule-based** grants were received under a legislation that grants \$600 (\$650 in the case of schools with Grades 8 and 9) to every school *irrespective of enrollment*. We treat this as 'anticipated' in our analysis. The second kind, **discretionary grants** were disbursed to schools at the discretion of the District Education Office as well as external donors. We treat this as unanticipated recognizing that this may overestimate the 'true' unanticipated component of cash grants .

[17] This was largely due to problems in the planning department of the Ministry of Education coupled with problems in procurement, rather than due to the lack of funds (less than 60% of the allotted budget was actually used during the fiscal year).

One concern is that legislated allocations may have little to do with received grants and our treatment of such grants as 'anticipated' may therefore be wrong.[18] A tracking exercise presents some encouraging results on this front (Table 1a, Cash-Grant Characteristics). In three out of the four provinces, over 93% of the schools surveyed received exactly the amount allocated. In the fourth province, Lusaka, this percentage dropped to 71% and this was attributed to delays in disbursement: Based on receipts in the previous year as well as interviews with head-teachers and district officials, it appeared that the remaining schools would receive the allocated disbursement shortly after the survey. It is precisely this delay that we exploit in drawing the distinction between w^a and $w^a_{received}$ in 12 above (details of the tracking exercise can be found in Das et al. 2002). Note also that the per pupil rule-based grants are fairly high in absolute amounts ranging from K 2,400 (urban) to K 8,700 (household sample), which corresponds to the cost of two textbooks or 36 boxes of chalk.

In contrast to the regularity of rule-based allocations, we find that the probability of receiving discretionary grants was much lower (24.15%). Conditional on receipts such funds tend to be either very small or extremely large- the inter-quartile range for log discretionary grants ranges between 6 and 10 log kwacha per pupil with a coefficient of variation greater than 6. This large dispersion in discretionary grants is also seen in Table 1a when we compare the logs to the actual amounts: In the case of actual amounts, discretionary are always larger than rule-based funds, in logs however, this relative ranking reverses due to the wide dispersion in the former. Finally, variables such as school or pupil characteristics have almost no predictive power in explaining the distribution of discretionary funds—less than .05% of the variation in such funds can be explained through differences in student composition, characteristics of teachers (head-teachers) or the availability of educational resources in the school.

The last point also addresses the potential targeting of discretionary funds (Table 1b). We find that, at least on the basis of observable outcomes, there is little difference between schools that received discretionary funds vis-a-vis those that did not. Although schools that received such funds tend to have students who are marginally wealthier and are located closer to the district office, these differences are not significant at the 10% level. Moreover, there is no difference in baseline scores between schools that received discretionary funds and those that did not—this result in particular, addresses the issues of mean reversion discussed earlier.

Thus, on the one hand rule-based allocations are clearly demarcated and defined, and schools receive the amount stipulated. On the other, discretionary funds are more volatile- less than a quarter of all schools receive such funds, and even conditional on receipts, the amount received varies dramatically. Further, such funds do not seem to have been allocated in a targeted fashion, at least on the basis of observable school and student characteristics. It is precisely this difference that forms the basis for our division cash grants into anticipated and unanticipated components.

[18] In the case of Uganda for instance, Ablo and Reinikka (2000) showed that less than 30% of the allocated capitation grant was received by schools.

5.4. Description: Households

To complete our description, it is also instructive to examine *household* funding of school inputs in comparison to the funding received from the government. In particular, if household educational expenditures are small compared to school grants, anticipated grants may play a role in the alleviation of credit-constraints at the household level (although our empirical test would remain valid). Figure 3 explores the importance of different funding sources for educational expenditure, disaggregated by schools that received high/low anticipated cash grants (with the cutoff at the median). We find that in both types of schools, household expenditure accounts for a large share of total (public and private) spending on education, ranging from 25% to 33% across schools that received high/low cash grants. The other significant expenditure share is accounted for by teachers salaries (roughly 50% in both cases); ignoring this component implies that household expenditure accounts for between 50% and 60% of total available expenditures. Since the household data is based on a sample of 'remote' schools that tend to be poorer, this percentage represents a *lower bound* on the actual share of household expenditure in total funding. Clearly then, even in remote and poor areas, households represent an extremely important component of educational funding and it is likely that they have sufficient leeway to adjust for changes in expected grants at the school level.

In anticipation of our results, Table 2 then looks at key household and school variables for schools with high/low enrolments (and hence high/low anticipated grants). The first row (matching success) shows the percentage of children in the primary age group who were successfully matched to the surveyed school, verifying our assumption of no school choice through the choice of the sample. For both high and low grant schools, matches are above 95%, but there is a (significant) 3% difference suggesting that endogenous enrolment may still be an issue.

The next two rows describe school cash grants and household expenditures. There are large differences in the means of two groups, although interestingly total funding is roughly equivalent, at K24,000 in low and K22,000 in high grant schools. The other rows in the table correspond to observable components of schools and households. For a number of important variables (household's asset indices, mother's/father's education and village enrollment) there is no significant difference between the two categories of cash grants. Moreover, in cases where differences are significant (percentage with mother/father at home) the direction is the opposite of what we would expect if enrolment was endogenous to villages and schools—high enrolment (low grant) schools tend to have *fewer* children with parents at home and report fewer visits from teachers to the household.

These statistics thus suggest that (a) there are significant differences in household contributions across low and high cash grant schools, although the null hypothesis that total funding is the same cannot be rejected (b) while there are some differences in observable household components across the two categories, these differences tend to be small or of the wrong sign. Our results on the substitutability of household and school cash grants verifies these broad results in a structured manner.

5.5. Description: Cognitive Achievement

The Examination Council of Zambia administered tests in Mathematics, English and the vernacular following the sampling scheme described above and the same children were tested one year apart. Although we should hence have a two year panel of test-scores for 3,500 children (since there were less than 20 children in Grade V in some schools), attrition in the data set allows for a smaller sample of 2,587 children. This drop is attributable to a number of factors including school-transfers/drop-outs (30%), absence on the day of the test (50%) and data issues arising from the inability to survey some schools or adequately complete pupil rosters (20%). We find some differences in original scores, with those who were unable to take the second test reporting significantly lower English and Math scores in the first year (0.11 and 0.19 s.d. respectively), but do not find any systematic pattern in attrition across schools receiving different amounts of anticipated/unanticipated cash grants.

Turning to learning gains over the year, students on average were able to answer 3.2 questions more in Mathematics from a starting point of 17.2 correct answers (from 45 questions) and 2.4 more in English starting from 11.1 correct answers (from 33 questions). In terms of our latent distribution, children gained on 0.42 standard deviations in Mathematics and 0.40 in English. These gains are disappointingly small- after one year of teaching students were able to increase their scores by 6% and 7.5% in Mathematics and English respectively.

Finally, Figure 4 shows the characteristics of the Mathematics and the English test with respect to the standard error of our latent score distribution. For both Mathematics and English, the test was 'too-hard' in the sense that children at the lower end of the distribution have (much) higher standard errors than those above the mean. Further, the English test was *better* designed than the Mathematics with lower estimated standard errors at all points of the distribution. Following from our previous discussion we thus expect considerable noise in our estimates, but also lower standard errors for our cognitive achievement results based on the English compared to the Mathematics test.

6. Results: Does Household Spending substitute for School cash grants?

Our main interest in this section is the estimation of the equations given by Equations 11,12 and **??**. To recapitulate, equation 11 estimates the relationship between household and school grants using anticipated grants that had already been received at the time of the survey ($w^a_{recieved}$). Equation 12 then uses the anticipated grants (w^a) instead (whether received or not at the time of the survey); if households truly anticipate such grants, the elasticity of substitution based on the second equation should be greater than the first. Finally, 13 then checks for the importance of credit constraints by examining differential responses across households categorized by wealth.

The estimation results are presented in Table 3a and 3b. For every estimated equation there are two specifications. The first is the Tobit specification where the sample includes all school going age children and educational expenditures on children who are not enrolled is defined to be zero. We may be concerned that the Tobit specification does not entirely capture the error structure of 11, with clustering at the village level. To account for this clustering, we also present estimates from a random effect Tobit specification which accounts for such clustering, but at the cost of a structural assumption on the clustered error terms. Columns (1) and (2) thus report results from 11, (3) and (4) from 12 and (5) and (6) from **??**. Table 3b then interprets these coefficients as the marginal impact (computed at the mean of the regressor) and the probability that the dependant variable is uncensored.

The results are as predicted and robust to the sample and specification used. Using $w^a_{received}$ as the regressor, the estimated elasticity of substitution for anticipated grants is always negative and significant ($\widehat{\beta_2}(w^a_{recieved}) < 0$) and ranges from -0.26 (Tobit) to -0.27 (Random-Effects Tobit). Further, $\widehat{\beta_2}(w^a) < \widehat{\beta_2}(w^a_{recieved})$ with the elasticity of substitution increasing to -0.49 (Tobit) to -0.50 (Random Effects Tobit) suggesting that households truly anticipate these cash grants and make their expenditure decisions prior to their actual receipt. Moreover, in 3 specifications the coefficient of unanticipated grants is small and insignificant ($\widehat{\beta_3} = 0$). For the specification where $\widehat{\beta_3} > 0$, the size of the coefficient is less than half that of $\widehat{\beta_2}(w^a_{recieved})$ suggesting that there may have been some household responses, even to unanticipated funding, but this was relatively small compared to funding that was anticipated. Note that to the extent that households *did* respond to unanticipated funding as well, this would imply that our coefficient on such funding in the test score regression is an underestimate of the true production function parameter.

Interestingly we always find the probability that the dependant variable is uncensored is lower for schools with higher anticipated funding (and therefore higher enrollment by construction). One interpretation is that this represents a quality-quantity trade-off whereby villages that have a lower preference for education also have higher fertility and our estimates thus represent a lower-bound for the true elasticity of household responses. Finally, columns (6) and (7) suggest that credit constraints may not be important (or that our asset index is not sufficiently accurate to discern this impact).

These results present strong evidence for a high elasticity of substitution between anticipated grants and household funding, and support the hypothesis that households make their educational funding decisions prior to the actual receipt of such grants. Further, households do not respond to unanticipated grants—despite the comparability of such grants (in amounts) to the anticipated equivalent, there is little evidence that households alter their own behavior as a response.

Our main worry with these results is the possibility of omitted variable bias that arises due to the close link between anticipated grants and enrollment and this is addressed in Table 4a and 4b. Following the strategy outlined in Section IV.2 we estimate the determinants of log anticipated/received funding in the first stage and use the residuals as an additional regressor in the specifications estimated under 11 and 12.

When we use received funding the residual is significant and positive, thus rejecting the test of weak exogeneity. This is counter-intuitive since if the unobserved village level attributes increases both enrollment and expenditure conditional on enrollment, we would expect the coefficient on the residuals to be *negative* instead (recall that our dependant variable is the inverse of enrollment). Returning to the difference between anticipated and received funding, note that by using received funding as the dependant variable we have *imposed* high enrollment for schools that received zero funding at the time of the survey in combination with a lower than expected expenditure. It is precisely this imposed relationship that drives the significant positive residual in column 1 and 2. Once we correct for this by using anticipated funding instead (Columns 3 and 4), the coefficient is insignificant at the 15% level confirming that the log of anticipated funding is exogenous to child level educational expenditures and estimated coefficients thus represent a causal relationship.

7. Results: Test Score Hypothesis

7.1. Graphical Evidence

The results in the previous section provide strong evidence that school grants and household funding are indeed substitutes—greater cash grants given to the school reduces the amount spent by the household. As discussed in Section IV.2, when school and household inputs are technical substitutes, the impact on learning gains of unanticipated is larger than that of anticipated inputs.

Figure 5 explores this relationship through non-parametric plots of the relationship between cash grants and gain in cognitive achievement. The figure on the left shows the (annual) change in cognitive achievement in the subjects of Mathematics and English plotted against (log) unanticipated grants while the relationship between gains in cognitive achievement and (log) anticipated grants. In both figures, the left axis shows the density of cash grants plotted on a histogram while the right axis depicts learning gains in Mathematics and English, plotted against the log grants. The figure verifies our basic hypothesis: learning gains are higher for unanticipated compared to anticipated grants. In the case of Mathematics there is a gain of almost 0.2 standard deviations and for English 0.15 standard deviations moving from the minimum to maximum unanticipated grants. Moreover, there is no discernible pattern in the case of Mathematics and a decline in the case of English for anticipated grants.

Figure 5 also suggests reasons for high standard errors in our estimation procedure. From the histogram for unanticipated grants it is clear that the distribution is marked by a large percentage of schools that receive zero combined with substantial variation among those that do receive positive amounts. Consequently, large variation in learning gains among the non-receivers might decrease the precision of the estimated relationship between cash grants and learning gains. Further, there appear to be differences in the precise functional form between Mathematics and English: For Mathematics small amounts of unanticipated grants have a low impact on learning achievement while for English such investments have an impact but decreasing returns set in quickly at higher levels. Although ideally we would like to estimate these non-linearities, our sample of 42 schools that receive positive amounts precludes further sub-division. Thus, we account for non-linearities through the inclusion of a quadratic term for unanticipated grants in equation 14 and check for robustness using a dummy variable for schools that received non-zero unanticipated grants.

7.2. OLS Results

Table 5 shows the results for English and Mathematics for four different specifications where all estimations are at the school level.[19] The first specification (column 1 for English and column 2 for Mathematics)

[19] Since the estimated equation is linear, averaging over children should have no impact on the estimated coefficients. In fact, child level regressions show similar coefficients but the signifncance is reduced when the regression is clustered at the school

is the most parsimonious and includes only anticipated and unanticipated grants in the estimation. The next two columns then include four additional explanatory variables- a dummy for rural is included to proxy for 'shifters' and three variables that capture potential changes in user–costs—whether the head-teacher has changed, whether the head of the parent-teacher association changed and the difference in PTA fees. The third specification (columns 7 and 8) include as variables the portion of anticipated funds that had actually been received by the time of the survey, as in Section V above.

For all specifications we find that the coefficient on anticipated grants is small and insignificant (except for the negative and significant value in column 7 for English), marginally more so in the case of English compared to Mathematics. For English the coefficient on the linear term of unanticipated grants is always positive and significant, and for the quadratic term negative and significant. For Mathematics the results are not as sharp; the coefficient on unanticipated grants is smaller and insignificant in 3 specifications, but is significant when unanticipated grants are treated as a dummy variable. Treating cash grants as a continuous variable, these results imply that the marginal impact of (log) grants on cognitive achievement at the mean is .048 standard deviations for English and .029 for Mathematics, representing 12.5% and 7% of the annual gain in learning. The results from the dummy specification imply that schools receiving non-zero unanticipated grants improve their learning by 0.12 (English) and 0.09 (Mathematics) standard deviations, which corresponds to approximately one-third of the average gain through the year.

There are two observations on the set of estimated coefficients. First, returning to the standard error of our latent score distribution, it is clear that the measurement error in the case of Mathematics is much larger than that for English. This would suggest that our estimates our more precise for the latter and could explain the significant findings for English but not Mathematics. Second, in Figure 5 the differences in English are primarily driven primarily by schools receiving zero and positive grants while those in Mathematics are driven by the difference between those receiving small and large amounts. In using the log specification, the optimal constant added to those with zero unanticipated grants is only K3.73. Given the steep gradient of the logarithm function near zero, the addition of a small constant implies that the estimated relationship is driven almost entirely the difference between the 'zero' and the 'non-zero' group rather than by the 'low positive' and the 'high positive' receivers. Consequently, in the case of Mathematics, the positive relationship between the low and the high group is overwhelmed by the flat portion before with the reverse for English. However, when we use a dummy variable for whether or not the school received any unanticipated funds as a regressor, both coefficients are significant and positive (and fairly close to each other) since we 'average' out these functional differences.

7.3. IV Results and Comparison

The final set of results we present corrects for the potential placement problem through the instrumentation strategy detailed above. The first stage from these results are presented in Table 6, and the results on learning changes using the instrumented unanticipated funds in Table 7. In the first stage, we find that the

rather than the district level. All coefficients retain their significance at least at the 10% level of confidence. Further, the results from the IV estimation are identical both in the size and the significance of the coefficients. These differences may arise due to the clustering of errors at both the school and the district level.

grants that the district received from external sources significantly impacts on the amount of unanticipated funds that a school will receive; the only other significant variable is whether the district was one that was heavily contested in the run-up to the election the same year as the survey. Together, the lagged variables and the funds received through external sources account for 50% of the variation in unanticipated funds at the level of the school, conditional on receipt. For the second part of the hurdle model, the probability that a school receives any cash grants is determined largely by the status of the school as one that was eligible to receive funds from an external donor program (the Program for the Advancement of Girls Education) administered by the United Nations since 2000.

Turning to the instrumented estimation results, we find that in all cases, there is an increase in the linear and a decrease in the quadratic term of the estimated equation. Moreover, for both English and Mathematics, the coefficients on unanticipated funds and unanticipated funds squared are always significant, while on anticipated funds the coefficients continue to remain small and insignificant. However, due to the quadratic formulation, it is not immediately apparent that these findings can be interpreted as evidence of 'positive-placement' whereby schools that were likely to do worse received more grants. Thus, to formally evaluate our hypothesis that the marginal impact of unanticipated is greater than that of anticipated grants, the following table presents results of Wald tests at various points of the sample range.

Subject and Estimation Strategy	Test Rejects Equality of Coefficients Evaluated at (p values)				
	Sample mean	25th %tile	Median	75th %tile	90th %tile
English (OLS)	5%	5%	5%	5%	>10%
English (IV)	1%	1%	1%	5%	5%
Math (OLS)	>10%	>10%	>10%	>10%	>10%
Math (IV)	10%	5%	10%	>10%	>10%

Since the marginal impact of unanticipated grants is given by $b_1 + 2b_2X$ where b_1 is the coefficient on unanticipated, b_2 on unanticipated grants squared and X the point at which the marginal value is evaluated, the Wald test takes into account the $var(b_1 + 2b_2X)$ given by the usual expansion. Every cell in the table shows the confidence level at which the null hypothesis that the marginal impact of unanticipated equals that of anticipated grants can be rejected. For English we find that the OLS results allow us to reject the null hypothesis at all points of the sample range at the 5% level, and we are unable to reject the null at the 90th percentile. The IV results are stronger; the null hypothesis can now be rejected at the 1% level for the mean, the 25th percentile and the median, and the 5% level for the 75th and 90th percentile. For Mathematics, overall results are weaker. We cannot reject the null hypothesis at any point of the sample range for the OLS results. With instrumental variables, results are marginally better with rejections at the 10% level for the sample mean and median, and at the 5% level for the 25th percentile.

8. Discussion and Policy Implications

Using data on learning achievement and non-salary cash grants to schools in Zambia we have tested a model of dynamic household optimization with the key prediction that anticipated compared to unanticipated grants will have a smaller effect on cognitive achievement if household and school cash grants are substitutes.

In the case of Zambia, we find that the elasticity of substitution between household and school grants ranges from -0.45 to -0.55. Consistent with our predictions, we then find that anticipated grants have zero impact on learning gains while unanticipated funds increases learning by 0.05 (English) and 0.025 (Mathematics) standard deviations at the mean. The estimated coefficients are likely under-estimates of the true production function coefficient due to the potential contamination of unanticipated grants by anticipated components. These results are robust to omitted variables arising either from school-choice or time-varying school effects. These results have implications both for estimation and policy and we discuss each of these in turn.

8.1. Implications for Estimation

The dominant estimation technique for estimating the effect of school inputs on cognitive achievement is based on the production function approach, where achievement (or changes in achievement) is regressed on such inputs. Following Todd and Wolpin (2003), these estimates back out the policy-effect of school inputs that combines both the effect of inputs on cognitive achievement through the production function as well as household responses to such inputs. Our use of unanticipated inputs allows the estimation of both effects separately, thus shedding more light on the process through which school inputs may or may not affect educational attainments.

For estimation purposes, it may appear that the problem in the production function approach arises due to omitted variables—if the researcher had access to both household and school expenditures in the current period, she would be able to accurately estimate the effect of the input through the production function. While true in a static setting, this does not take into account the possibility of inter-temporal substitution in a dynamic problem. Specifically, households start responding to school inputs at the time that information regarding such inputs is revealed so that the entire history of household inputs will be required from that point onward to estimate unbiased production function parameters. The alternative approach followed in this paper, is to examine the portion of inputs that arrive as exogenous shocks so that households are unable to respond in the current (or preceding) periods. In the absence of data on historical household inputs (as well as details on the revelation of information) the use of unanticipated inputs allows identification of the production function parameter with greater accuracy.

The distinction we draw between anticipated and unanticipated inputs could also account for the wide variation in estimated coefficients of school inputs on cognitive achievement (Glewwe, 2002, Hanushek, 2003 or Kreuger, 2003) The production function framework does not separate anticipated from unanticipated inputs so that the regressor is a combination of these two different variables. Since the covariance between the two types of inputs is (necessarily) zero, the estimated coefficient is bounded below by the 'policy effect' and above by the production function parameter; the distance from either bound depends on the extent to which the schooling inputs were anticipated or not.

There are a number of extensions that can be pursued in future research. While the use of the Euler equation framework allows us to control for heterogeneity arising from a number of sources and in a fairly flexible form, our model does require some restrictive assumptions. First, we are unable to control for heterogeneity due to the underlying production function or the household discount rate. Second, we are unable to allow for non-linear effects of the lagged test score and third, our specification requires the use of an additive lagged test score in the production function. The last two problems could potentially be addressed by panels that span a larger time period.

8.2. Implications for Policy

The argument developed here also has repercussions for educational policy. Our results *do not* suggest an educational policy where inputs are provided unexpectedly. Although cognitive achievement in the current period does increase with unanticipated inputs, the additional consumption now will push them off the optimal path. In subsequent periods, therefore, they will readjust expenditures until the first order conditions are valid again—unanticipated inputs in the current period will not have persistent effects in the future (except due to the durable nature of the good).

The policy framework that *is* suggested under this approach involves constructing appropriate 'spheres of influence'. Under such a framework schooling inputs are characterized by the level of market imperfections that characterize their provision, as well as the degree of substitutability with household contributions. Inputs are then divided into two categories- those within the sphere of influence of either the government or the household.

The former would include inputs that are characterized by incomplete markets or other market failures (for instance, see Miguel and Kremer, 2003 for an example of market failures due to externalities) while the latter would consist of inputs provided under competitive market conditions. These would include inputs whose provision in the private market is characterized by market imperfections. An important example of inputs in the governments sphere of influence maybe teaching inputs, whereby problems in the design of contracts (see for instance Holmstrom and Milgrom, 1991 for problems arising due to multi-tasking or Holmstrom, 1982 for problems arising due to moral hazard in teams) may make public more efficient than private provision. Similarly, inputs that retain some aspects of public-goods and would thus be under-provided in the private market are a good candidate for government provision. Interestingly, both of these have been shown to have significant impacts on cognitive achievement (see for example, Hanushek (2001) on the importance of teachers and Jacoby (1994) on infrastructure).

To the extent that the government may be worried about equity in the provision of education, direct subsidies to households for inputs of the second kind characterize the optimal policy. The approach followed here of treating cognitive achievement as a household maximization problem with the production function acting as a constraint opens up a new avenue for research, one that explicitly recognizes the centrality of households and classifies schooling inputs into the categories discussed above. To construct the appropriate spheres-of-influence it would be important to identify characterize inputs by their degree of substitutability vis-a-vis household provision. To do so requires both the use of matched household and school surveys as well as the careful identification of surprises in the provision of inputs; in particular, long-term data on schooling inputs would allows for cleaner identification based on deviations from means, much as in the consumption literature.

References

Ablo, Emmanuel and R.Reinikka (2000) “Do Budgets Really Matter: Evidence from Public Spending on Education and Health in Uganda” Policy Research Working Paper 1926. The World Bank, Washington

Alderman, H, J.R. Behrman, S.Khan, D.R.Ross and R.Sabot (1997) "The Income Gap in Cognitive Skills in Rural Pakistan". Economic Development and Cultural Change, Vol. 46(1), pp.97-123.

Angrist, J.D. and V.Lavy (1999) "Using Maimonides' Rule to Estimate the Effect of Class Size on Student Achievement". Quarterly Journal of Economics, Vol. 114(2), pp.533-575.

Banerjee, A. (2003) “Educational Policy and the Economics of the Family”. Mimeo. Massachusetts Institute of Technology.

Banerjee, A., S.Cole, E.Duflo and L.Linden (2003) "Remedying Education: Evidence from Two Randomized Experiments in India". Mimeo. Massachusetts Institute of Technology.

Becker, G. and N.Tomes (1986) "Human Capital and the Rise and Fall of Families" Journal of Labor Economics, Vol. 4(3), pp.1-39.

Birnbaum, Allan (1967), *"Some Latent Trait Models and Their Use in Inferring an Examinee's Ability"* in Lord, Frederic M. and M.R. Novick "Statistical Theories of Mental Test Score". Addision-Wesley Publishing Company. London.

Bock, R.D. and Aitkin, M. (1981). *"Marginal Maximum Likelihood Estimation of Item Parameters: An Application of an EM Algorithm". Psychometrica 46, 443-459.*

Bock, R.D and Lieberman, M. (1970) "Fitting a Response Model for n Dichotomously Score Items". Psychometrika 35, 179-197.

Blundell, R.W. and R. Smith (1986) "An Exogeneity Test for a Simultaneous Equation Tobit Model with an Application to Labor Supply", Econometrica, 54, 679-685.

Brown, P. (2003) "Parental Education & Child Learning: Human Capital Investments in Time and Money". Mimeo.Michigan University.

Case, A. and Deaton, A. (1999). 'School inputs and educational outcomes in South Africa', Quarterly Journal of Economics, vol. 114 (3) (August), pp. F1047-F84.

Cox, D. and E.Jimemez (1995) "Public Income Redistribution in the Philippines" in Van de Walle and Nead (eds.).

Das, Dercon, Habyarimana and Krishnan (2002). "Rules vs. Discretion: Public and Private Funding in Zambian Education". Mimeo. The World Bank.

Deaton, A. and Muellbauer, J. (1980) "Economics and Consumer Behavior". Cambridge University Press.

Dercon, S. and P. Krishnan (2000) "In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia", Journal of Political Economy, vo. 108 (4), pp.688-728.

Filmer, D., and L.Pritchett (1999) “The Effect of Household Wealth on Educational Attainment: Evidence from 35 Countries”. Population and Development Review 25(1): pp.85-120.

Foster, A.D. (1995) "Prices, Credit Markets and Child Growth in Low-Income Rural Areas" The Economic Journal, Vol. 105, No. 430, pp. 551-570.

Glewwe, P. (2002) "Schools and Skills in Developing Countries: Education policies and socioeconomic outcomes" The Journal of Economic Literature,. Vol XL, pp.436-482.

Glewwe, P. and H.Jacoby (1994) "Student Achievement and Schooling Choice in Low-Income Countries: Evidence from Ghana", Journal of Human Resources, Vol. 29 (3) pp.843-864.

Greene, W.H.(2002) "Econometric Analysis"

Hanushek, E.A. (1971). 'Teacher characteristics and gains in student achievement: Estimation using micro data', American Economic Review, vol. 60 (2) (May), pp. F280-F8

Hanushek, E.A. (1986). 'The economics of schooling: production and efficiency in public schools', Journal of Economic Literature, vol. 24 (3) (September), pp. F1141-F77.

Hanushek, E.A. (1997) "Assessing the Effects of School Resources on Student Achievement: An Update". Educational Evaluation and Policy Analysis. Vol. 19(2), pp. 141-164.

Hanushek, E. A. (2003). 'The failure of input-based schooling policies', Economic Journal, vol. 113 (February), pp. FF64-F98.

Holmstrom, B. (1982) Moral hazard in teams. Bell Journal of Economics, Autumn.

Holmstrom, B. and P.Milgrom (1991) Multitask principal-agent analysis: Incentive contracts, asset ownership and job design. Journal of Law, Economics and Organization 7: 24-51.

Hoxby, C. (2000) "Peer Effects in the Classroom: Learning from Gender and Race Variation". NBER Working Paper #7867, Cambridge, Massachusetts.

Jacoby, Hanan G., “Is there an Intrahousehold ‘Flypaper Effect’? Evidence from a School Feeding Programme.” *The Economic Journal* 112, 196-221, January 2002.

Jacoby, H.G. and E. Skoufias “Risk, Financial Markets, and Human Capital in a Developing Country,” Review of Economic Studies, Vol. 64, No. 3, July 1997, pp. 311-335

Johnson, S.R and G.C.Rausser (1971) "Effects of Misspecifications of Log-Linear Functions when Sample Values are Zero or Negative" American Journal of Agricultural Economics, vol. 53 (1), pp.120-124.

Kane, T.J. and D.O.Staiger (2001). "Improving school accountability measures". Working paper #8156, National Bureau of Economic Research. Cambridge, MA.

Kane, T.J. and D.O.Staiger (2002). "The promise and pitfalls of using imprecise school accountability measures". Journal of Economic Perspectives, 16(4), pp.91-114.

Krueger, A. (2003).'Economic considerations and class size', Economic Journal, vol. 113 (February), pp. FF34-F63

Lazear, E. P. (2001). 'Educational production', Quarterly Journal of Economics, vol. 116(3), pp. 777-803

Miguel, Edward, and Michael Kremer, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities". *Forthcoming Econometrica,* 2003.

Moffit, R. (1992). "Incentive Effects of the U.S. Welfare System: A Review". Journal of Economic Literature. Vol. 30, pp. 1-60.

Murphy, Kevin, and Robert Topel, (1985) "Estimation and Inference in Two-Step Econometric Models," Journal of Business and Economic Statistics, III (1985), 370-79

Ravallion, M. and G.Datt (1995) "Is Targeting Through a Work Requirement Efficient? Some Evidence for Rural India" in Van de Walle and Nead (eds)

Rosenzweig, M. and K. Wolpin (1994) "Parental and Public Transfers to Young Women and Their Children". American Economic Review. Vol. 84, pp. 1195-1212.

Todd, P.E. and Wolpin, K.I. (2003). 'On the specification and estimation of the production function for cognitive achievement', Economic Journal, vol. 113, pp. F3-33

UNESCO

Urquiola, M.(2001) "Identifying class size effects in developing countries: Evidence from rural schools in Bolivia". World Bank Working Paper Series, Number 2711.

Van de Walle, D. and K.Nead (eds) 1995 "Public Spending and the Poor: Theory and Evidence" Baltimore, Johns Hopkins University Press for the World Bank.

Wooldridge, J.M. (2001) "Econometric Analysis of Cross Section and Panel Data". MIT Press.

World Bank

9. Appendix 1

Proof of 4. To define the user cost, we consider a change in z_t and z_{t+1} that increases cognitive achievement by one unit in period t only. Formally,

$$dTS_t = F_{z_t}(w_t, z_t, \mu, \eta)dz_t \tag{17}$$

$$dTS_{t+1} = F_{TS_t}(w_t, z_t, \mu, \eta)dTS_t + F_{z_{t+1}}(w_{t+1}, z_{t+1}, \mu, \eta)dz_{t+1} = 0 \tag{18}$$

where F_{z_t}, $F_{z_{t+1}}$ and F_{TS_t} are marginal effects of the educational achievement production function with respect to household inputs in the current and next period, related to initial levels of cognitive achievement in the next period. The conditioning (w_t, z_t, μ, η) explicitly recognizes that these marginals will be a function of current levels of inputs as well as child and school non-time varying characteristics. To achieve this, assets are reduced in period $t+2$ by:

$$\begin{aligned} dA_{t+2} &= (1+r).(dz_{t+1} + (1+r)dz_t) \\ &= (1+r)(\frac{-F_{TS_t}(.)}{F_{z_{t+1}}(.)} + \frac{(1+r)}{F_{z_t}(.)})dTS_t \end{aligned} \tag{19}$$

which defines the user cost in the expression given by 4.

Maximum Likelihood Derivation of the latent variable θ Following Birnbaum (1967), we follow a parametric approach and use the 3 parameter logistic to map the latent variable θ to the probability of a correct answer in question j, $P_j(\theta)$ so that

$$P_j(\theta) = c_j + (1-c_j)\frac{1}{1+\exp\{-a_j(\theta - b_j)\}} \tag{20}$$

Following IRT terminology, the parameter b_j measures the difficulty of the item (a location parameter), a_j measures the discrimination of the item and c_j measures the guessing probability. We can then define the likelihood function as follows

$$L(\theta, a, b, c) = \prod_j \prod_i P_j(\theta_i, a_j, b_j, c_j)^{x_{ij}} \{1 - P_j(\theta_i, a_j, b_j, c_j)\}^{1-x_{ij}} \tag{21}$$

where x_{ij} is the response by individual i to question j. Maximization of the likelihood function then gives us the required normal equations. We use a Marginal Maximum Likelihood estimation (Bock and Lieberman, 1971) in combination with an Expected-Maximization (EM) algorithm (Bock and Aitkin, 1980). Under this scheme, a density function is assumed for the latent variable, $f(\theta)$. This is then 'integrated-out' to arrive at consistent estimates of the item parameters (a_j, b_j, c_j). As the last step the item parameters are then used to calculate θ. Finally, Fisher's information measure for the latent variable θ, can be written as

$$\begin{aligned} I(\theta) &= \sum_{j=1}^{J} I_j(\theta) \text{ where} \\ I_j(\theta) &= \frac{\{P_j^{'}(\theta)\}^2}{P_j(\theta)(1-P_j(\theta)} \end{aligned} \tag{22}$$

The standard error or our estimate $\widehat{\theta}$ is now simply $\frac{1}{\sqrt{\sum_j I_j(\theta)}}$ and by the property of ML estimators $\widehat{\theta} \sim N(\theta, \frac{1}{\sqrt{\sum_j I_j(\theta)}})$

Figure 1: Determination of Remote Locations (Screen Shot of Northern Province)

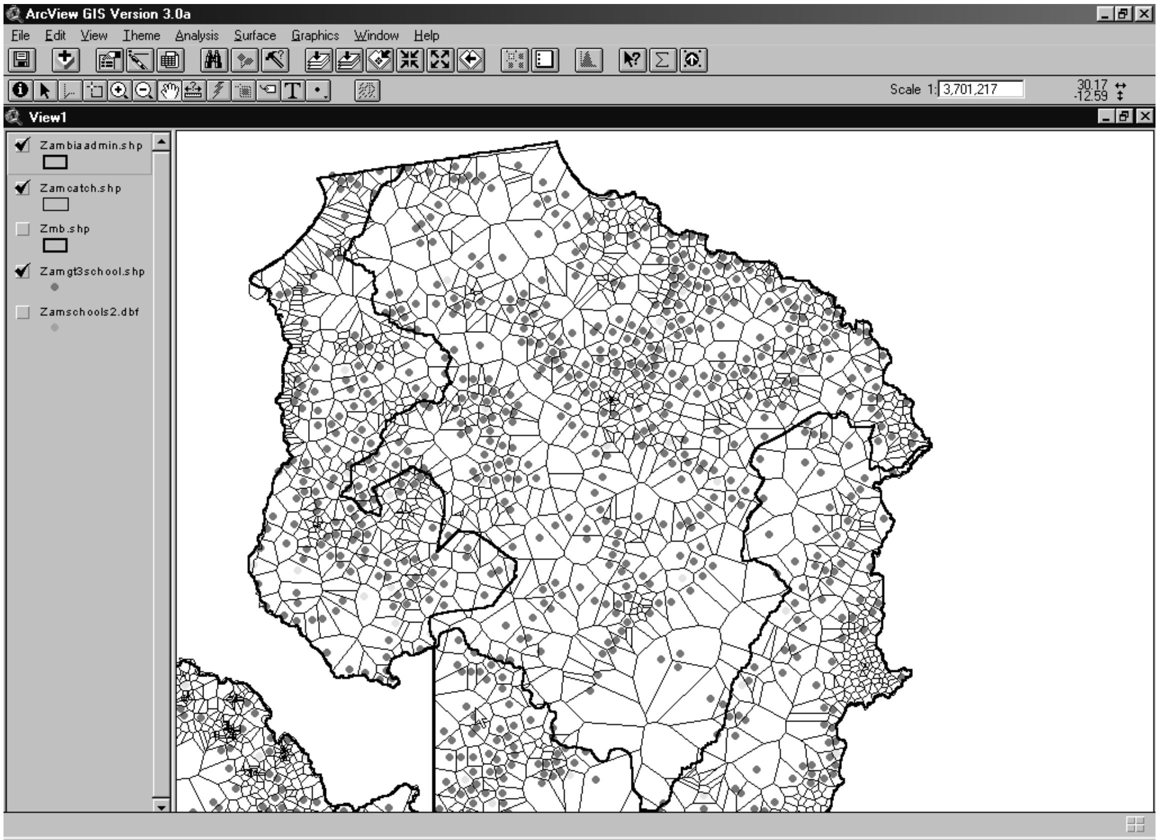

Note: The figure shows how schools were chosen for the household sample. Every point is a school and lines represent the edges of a Thiessen Polygon. On any line, every point is equidistant between two neighboring schools. Remote schools are those where the distance of the closest line to the school is at least 2.5 kilometers. Choosing only among those schools that are in the sample imposes further restrictions---these are shaded in yellow. The province shown in this figure is the Northern Province.

Figure 2a: Net Enrollment Rates in Africa

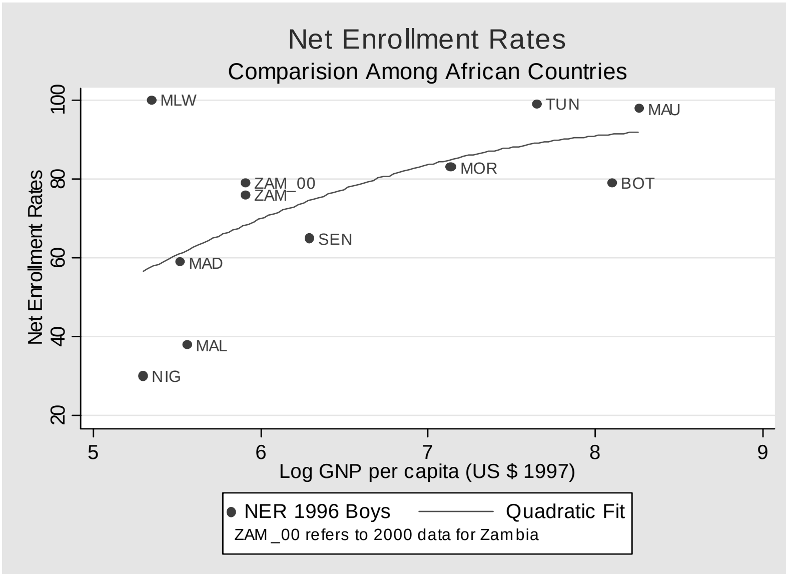

Source: UNESCO (2002). This graph shows the relationship between net enrollment rates and log GNP per capita in selected African countries, with a fitted quadratic. Zambia lies above the fitted line, suggesting that the enrollment in Zambia is great4er than what would be predicted through per capita income alone.

Figure 2b: Educational Attainment Curves: Zambia

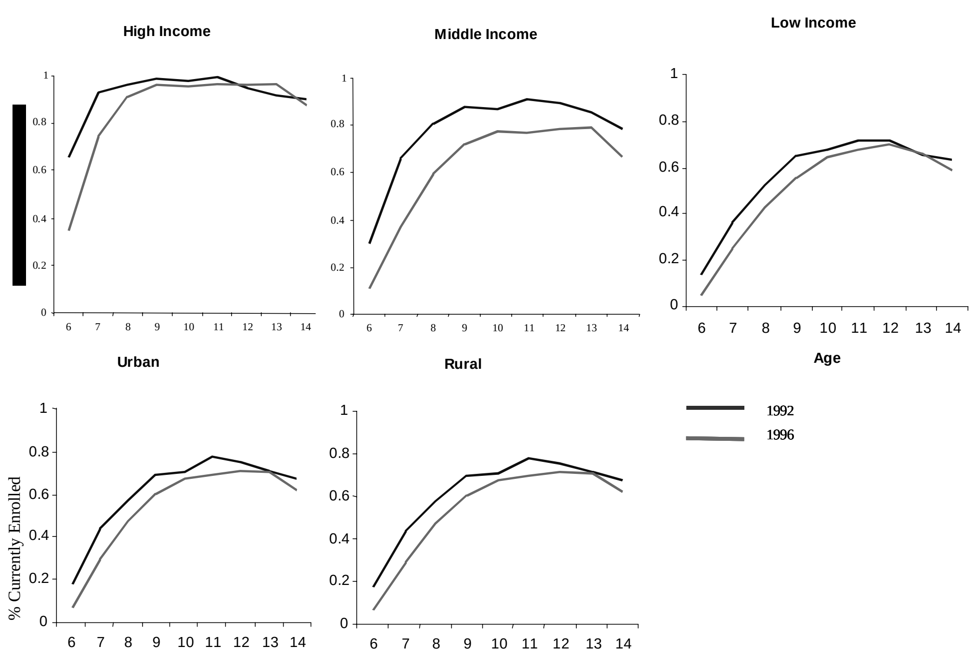

Source: Based on Filmer and Pritchett (1999)

These graphs show educational attainment plots of the percentage of children enrolled against age. For all regions and income categories, we find that strong evidence of delayed enrollment with enrollment increasing with age till age 12 and then tapering off or declining. The graphs also show how educational outcomes have worsened between 1992 and 1996, with a decline in enrollment at every age group and for all socioeconomic levels.

Table 1a: Descriptive Statistics: School Enrollment and Staffing

Category	Variable	Urban	Rural	`Remote' Schools (HH Sample)
Basic Indicators	School size (number of pupils)	1439.5 (600.98)	553.4 (408.19)	399.1 (224.3)
	Pupil-teacher ratio	42.23 (24.16)	63.62 (55.14)	66.1 (40.91)
	Number of pupils per classroom in good condition	103.4 (58.41)	96.7 (46.40)	101.4 (59.66)
Outcome Indicators	Repeating the same grade (%)	4.9 (4.29)	9.42 (6.66)	9.1 (5.33)
	Dropouts as ratio of current enrolment (%)	1.67 (2.57)	4.49 (5.05)	4.6 (5.28)
	Pass-rate in 1999 Grade VII examination (Males) [1]	40.5 (22.68)	44.2 (27.07)	42.6 (27.47)
	Pass-rate in 1999 Grade VII examination (Females) [1]	38.6 (24.32)	40.7 (30.73)	38.4 (30.24)
Pupil Characteristics	Average value of wealth-index of households with children in the school[2]	0.57 (0.61)	-0.56 (.56)	-0.73 (0.43)
	Percentage of children who are orphans	4.7 (3.6)	4.79 (4.13)	4.9 (3.5)
Cash-Grant Characteristics	Percentage of schools who received **anticipated funds** at time of survey	89.4 (3.7)	89.3 (2.9)	85.7 (5.9)
	Percentage of schools who received **unanticipated funds**	23.1 (5.2)	24.8 (4.0)	14.2 (5.9)
	Anticipated amount received (log Kwacha per pupil)	7.66 (0.41)	8.67 (0.60)	8.93 (0.54)
	Unanticipated amount (log Kwacha per pupil)	7.22 (2.31)	7.93 (2.58)	9.84 (2.77)
	Anticipated amount (Kwacha per pupil)	2372.8 (1535.4)	6931.4 (3969.7)	8676.1 (4852.1)
	Unanticipated amounts (Kwacha per pupil)	17121.2 (36822.5)	54559.0 (14424.7)	120630.9 (158621.6)

Source: ESD Sample. Standard-Deviations in brackets.

1. Pass-Rates are for the Grade VII examination administered to all students by the Examination Council of Zambia.
2. The wealth index is based on a weighted aggregation of household assets similar to a principal components analysis, but with weights optimally derived to minimize classification errors. Details are in "Discretion vs. Rules, Appendix 1".
3. For anticipated funding, 2.25% of schools received more than the allotted amount, to a maximum of $800.
4. Unconditional logs for cash-grants are calculated by generating lnX=ln(X+b) if X=0, where b is determined optimally .

Table 1b: Who Received Discretionary Funds?

Category	Variable	No Discretionary Funds Received	Received Discretionary Funds	Significant Difference?
School Characteristics	School Asset Index	-0.03 (0.79)	.09 (.77)	NO
	Total Enrollment in School	862	981	NO
	Distance to District Office (% within 5 KM)	48.1%	60.4%	NO
	Distance to Provincial Office (% within 5 KM)	20.7%	18.6%	NO
Performance in 2001 Examinations	English Scores in 2001	-0.045 (.48)	-0.05 (.52)	NO
	Mathematics Scores in 2001	-0.012 (0.47)	-.069 (0.42)	NO

Source: ESD Sample. Standard-Deviations in brackets. This table checks to see whether schools that received discretionary funds were `different' along observable dimensions from those that did not. We find no difference in either school characteristics or test-scores in 2001 between schools that received such funds versus those that did not. The wealth index is based on a weighted aggregation of household assets similar to a principal components analysis, but with weights optimally derived to minimize classification errors. Details are in Appendix I of Das et al. 2003. Test Scores are the maximum likelihood estimates of the latent variable as described in Appendix 1
.

Figure 3: Household Expenditures and School Funding

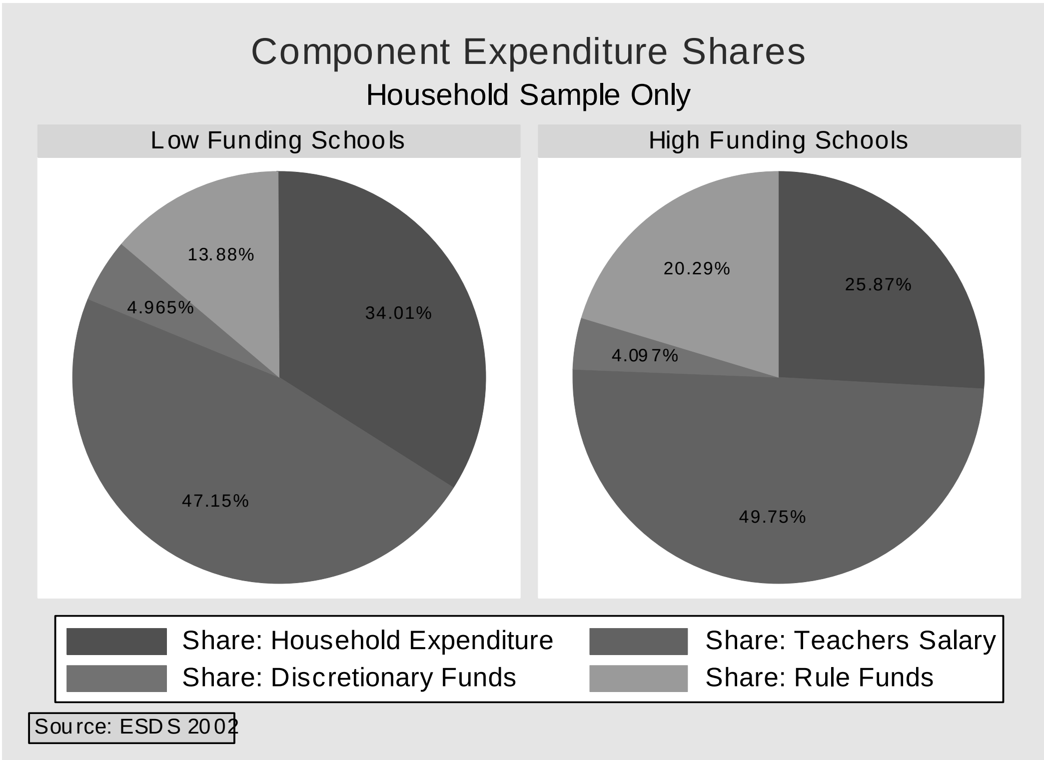

Source: ESD Sample. The pi-chart shows how educational inputs are funded in schools that received high/low anticipated funds. The shares are computed as the average of shares across schools.

1. Teachers Salary is computed as salary divided by the number of students in the teachers class. This is computed for a sample of teachers who were interviewed if they were either currently teaching Grade VI or Grade V students, or had taught Grade V students in the previous year. The particular sample was chosen to ensure that teacher characteristics could be matched to students who were tested in both years. Salaries will therefore be biased if there is selection of teachers into different grade levels.
2. Household expenditure is based on a one-year recall question of household educational expenditure for every child on various items including textbooks, school supplies and uniforms.
3. Discretionary funds are unanticipated by households.

Table 2: Cash-Grants and School Characteristics in High and Low Anticipated Grant Villages

Categories	Variable	Low Grant Schools	High Grant Schools	Significant?
Matching Success	% School going Children attending surveyed school	94.6%	98.5%	YES
Outcome Variable	Average Per-Child Expenditure	19,576	10,794	YES
Explanatory Variable (Anticipated Grants= K/enrollment)	Average Anticipated Grant	4,648	10,999	YES
	Total Enrollment in School	572.15	232.77	YES
Observable Components of Households	Schools Asset Index	-.75	-.93	YES
	Households Asset index	0.018	-0.054	NO
	% Children with mothers in household	72.5	79.8	YES
	% With fathers in household	65.0	72.5	YES
	% whose mothers can read	87.7	90.1	NO
	% whose fathers can read	72.5	79.8	NO
Observable components of schools	% who say Head-Teacher is good	0.79	0.84	NO
	% who say that Teacher visited household	52.9	74.2	YES
	%For whom school is within 30 minutes	50.6	70.7	YES
Village Enrollment	**% Enrollment in Village**	**.78**	**.80**	**NO**

Note: 1. Private Schools are not included, reducing the sample to 33 schools.

1. All tests of percentages are probability tests, all tests of continuous variables are t-tests.
2. Significant differences are at the 1% confidence level.
3. Schools Asset index is the average value of the asset index for children in the school matched to the household.

This table compares observable characteristics of households and schools for schools that received low and high anticipated funds respectively. We find that for a number of characterisitics, there is no diffference between the two categories. For variables where differences are significant (% with fathers in households or distance to school) the relationship with enrolment is the opposite of what one might expect, suggesting that there is no systematic correlation between enrollment at the school level, which determines anticipated funding, and household characteristics.

Figure 4: Test Characteristics English and Math

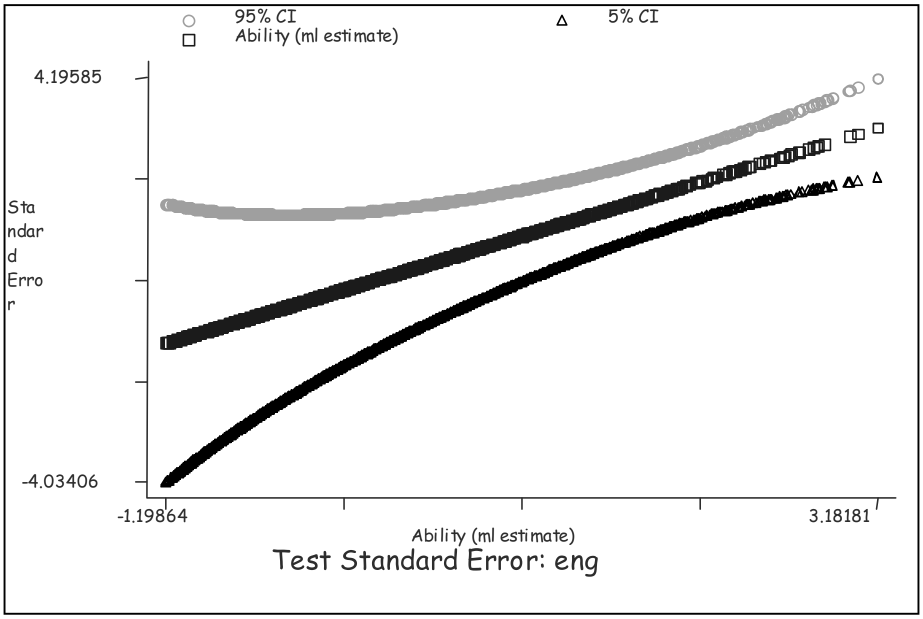

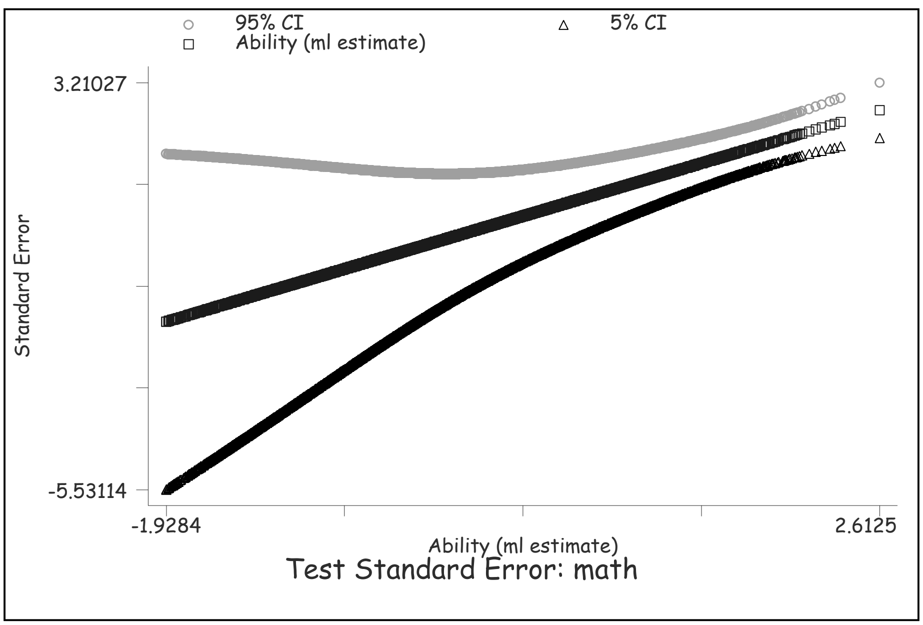

Source: Examination Council of Zambia data. These two graphs show the standard error of the estimation for the latent ability (knowledge) variable. The line in the middle plots ability against itself, while the lines on the outsides plot the 95% confidence intervals where the standard errors are calculated from Fisher's information for maximum likelihood estimates. Thus for instance, in both the Mathematics and the English examination, the exam was `too hard' so that the confidence band at lower ability levels is larger than that at higher ability levels. Further, the Mathematics exam is more `noisy' than the English exam with larger confidence bands, especially at the lower levels.

Table 3a: Relationship between Household and School Funding

	(1)	(2)	(3)	(4)	(5)	(6)
	Base Regression: Tobit Specification	Base Regression: Tobit with Random Effects	Hypothesis 2: Tobit Specification	Hypothesis 2: Tobit with Random Effects	Log Specification and Interaction with Assets	Tobit Specification: Asset Interactions
Log Received (rule-based) Funding	**-0.356 (2.51)***	**-0.338 (1.73)**				
Log Anticipated Funding			**-0.646 (4.95)****	**-0.657 (4.27)****	**-0.964 (3.75)****	**-0.541 (2.35)***
Log Discretionary Funding	-0.140 (2.31)*	-0.130 (1.58)	-0.072 (1.34)	-0.071 (1.13)	0.045 (0.45)	-0.089 (1.62)
Child is Male	-0.038 (0.32)	-0.039 (0.33)	-0.039 (0.33)	-0.040 (0.33)	-0.022 (0.28)	-0.023 (0.19)
Age	0.266 (13.03)**	0.268 (13.19)**	0.267 (13.18)**	0.267 (13.23)**	0.071 (4.21)**	0.271 (13.29)**
Household Wealth Index	0.730 (10.68)**	0.723 (10.06)**	0.731 (10.80)**	0.725 (10.36)**		
Medium Wealth Household					-3.729 (1.48)	0.176 (0.06)
High Wealth Household					-2.305 (0.62)	3.967 (1.42)
Medium wealth * Anticipated Funding					0.495 (1.78)	0.081 (0.25)
High Wealth * Anticipated Funding					0.382 (0.91)	-0.281 (0.91)
Constant	8.817 (5.98)**	8.573 (4.28)**	11.097 (8.95)**	11.187 (7.69)**	15.522 (6.11)**	9.381 (4.44)**
Observations	1407	1407	1407	1407	1197	1407
Log Likelihood	-2809.7	-2805.7	-2800.7	-2799.7	R2=0.18	-2807.8

Notes: The regressions in this table show the effect of anticipated and unanticipated funding on children's educational expenditures (the dependant variable in all regressions). Estimates marked ** are significant at 1%, * denotes significance at 5% Anticipated funding is further divided into funding that had been received by the school at the time of the survey, and funding that had not yet arrived (but did so later). Estimates from the former are presented in columns (1) and (2) and the latter in columns (3) and (4). All regressions use a Tobit specification with censoring at 0, the random effects Tobit specifications account for the clustering of errors at the level of the village. Marginal effects (conditional on being uncensored) and the probability of censoring are presented in Table 3b below. Further comments are (1) For all regressions, K100 is added on to discretionary funding to allow logs. The minimum funding is K900 conditional on receipt; (2) Province Dummies are included in all regressions; (3) The omitted category in (5) & (6) are the low-asset households and (4) The asset index based on optimal maximum likelihood weights. Results are robust to alternative indices (for instance an unweighted raw sum).

Table 3b: Marginal Effects and Probability of Censoring

	(1)		(2)		(3)		(4)	
	Base Regression: Tobit Specification		Base Regression: Tobit with Random Effects		Hypothesis 2: Tobit Specification		Hypothesis 2: Tobit with Random Effects	
	Marginal Effect at Mean	Prob. (Uncensored)	Marginal Effect at Mean	Prob. (Uncensored)	Marginal Effect at Mean	Prob. (Uncensored)	Marginal Effect at Mean	Prob. (Uncensored)
Log Received (rule-based) Funding	-0.27*	-0.02*	-0.26	-0.02				
Log Anticipated Funding					-0.49**	-0.039**	-0.50**	-0.04**
Log Discretionary Funding	-0.10*	-0.008*	-0.10	-0.008	-0.05	-0.004	-0.05	-0.004

Note: This table shows the marginal effects at mean values of the regressors based on the coefficients from Table 3a above. In all cases, tthere is no difference in using either the Tobit or the Tobit random effect specification, either in the size or the significance of estimated coefficients. The estimated elasticity increases substantially when we use anticipated instead of received funding, confirming that households make their educational expenditure decisions before such funding is actually received. Moreover, the probability that educational expenditures are uncensored (i.e. that children are enrolled) is *lower* in schools with higher anticipated funding. This implies that the probability of enrollment of any specific child is lower in schools with smaller enrollments, suggesting that any omitted variable bias leads to *lower* spending for enrolled children, reflecting perhaps a quality-quantity trade-off.

Table 4a: Test of Weak Exogeneity: Household Expenditures

	(1)	(2)	(3)	(4)
	Base Regression: Tobit Specification	Base Regression: Tobit Specification with Random Effects	Tobit Specification: Hypothesis 2	Hypothesis 2: Tobit Specification with Random Effects
Log Received (rule-based) Funding	-0.871 (3.96)**	-0.863 (3.11)**		
Log Anticipated Funding			-0.987 (3.71)**	-0.987 (3.61)**
Log Discretionary Funding	-0.154 (2.54)*	-0.148 (1.95)	-0.076 (1.41)	-0.076 (1.37)
Residuals from first stage	0.607 (3.06)**	0.614 (2.46)*	0.446 (1.47)	0.443 (1.42)
Constant	13.429 (6.38)**	13.295 (5.02)**	14.205 (5.80)**	14.204 (5.65)**
Observations	1407	1407	1407	1407

Note: The regressions presented here check for the weak exogeneity of anticipted funds using the approach followed by Blundell and Smith (1986). We estimate the determinants of anticipated funding (by proxy, enrrollment) in the first stage and re-estimate the specifications discussed in Table 3a with the estimated residuals. We reject weak exogneity of the actually received anticipated funding, but cannot reject that anticipated funds are exogenous to children's educational expenditures. The failure to reject in the former case arises from the imposed high enrollment when we use received funds in the first stage. All regressions reported here control for the same set of variables as those in Table 3a, specifically the gender and age of the child, the wealth of the household and province level dummies. As usual, ** denotes significant at 1%, * denotes significance at 5%

Table 4b: Marginal Effects and Probability of Censoring

	(1)		(2)		(3)		(4)	
	Base Regression: Tobit Specification		Base Regression: Tobit with Random Effects		Hypothesis 2: Tobit Specification		Hypothesis 2: Tobit with Random Effects	
	Marginal Effect at Mean	Prob. (Uncensored)	Marginal Effect at Mean	Prob. (Uncensored)	Marginal Effect at Mean	Prob. (Uncensored)	Marginal Effect at Mean	Prob. (Uncensored)
Log Received (rule-based) Funding	-0.66**	-0.05*	-0.66**	-0.05*				
Log Anticipated Funding					-0.75**	-0.06**	-0.75**	-0.06**
Log Discretionary Funding	-0.11*	-0.01*	-0.11*	-0.01*	-0.05	-0.004	-0.05	-0.004

Note: This table shows the marginal effects at mean values of the regressors based on the coefficients from Table 4a above. In all cases, there is no difference in using either the Tobit or the Tobit random effect specification, either in the size or the significance of estimated coefficients. The estimated elasticity is higher than that in Table 3b; this is consistent with substitutability between school enrollment and expenditures on children.

Figure 5: School Funding and Learning Gains

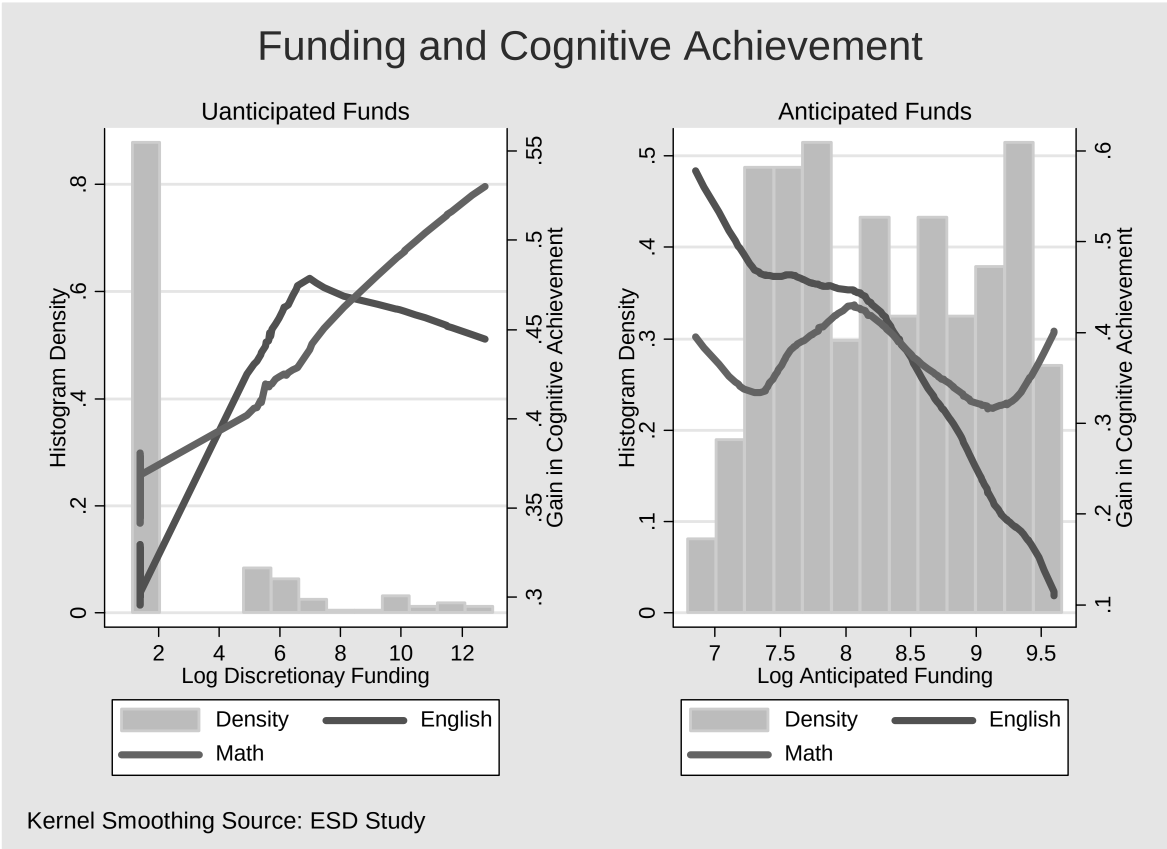

Notes: This figure shows the relationship between anticipated/unanticipated funds and gains in cognitive achivement. For both graphs, the histogram of funding is plotted on the left axis and the relationship between cognitive gains and funding is plotted on the right axis. In the case of unanticipated funds, a large number of schools receive 0 funds and conditional on receipt, there is very high variance in the amount received. For anticipated funds, the distribution mirrors the distribution of enrollment and is evenly distributed on the sample range. The relationship between cognitive gains and unanticipated funding is positive and significant for both Mathematics and English and not significantly different from zero for anticipated funding.

Table 5: Funding and Test Scores

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Basic Regression: English	Basic Regression: math	Full Control Set: English	Full Control Set: Math	Dummy on Discretionary: English	Dummy on Discretionary: Math	Full Control Set: Expected Rule: English	Full Control Set, Expected Rule: Math
Log of Unanticipated Grants/ Dummy	0.076 (2.54)*	0.027 (1.22)	0.072 (2.42)*	0.028 (1.42)	0.125 (2.28)* (coefficient for dummy on discretionary)	0.095 (2.07)* (coefficient for dummy on discretionary)	0.055 (1.92)	0.032 (1.37)
Square of log Unanticipated Grants	-0.006 (2.35)*	-0.001 (0.67)	-0.005 (2.06)*	-0.001 (0.65)			-0.004 (1.56)	-0.002 (0.86)
Log of anticipated Grants/ Dummy (received)	-0.025 (1.27)	0.007 (0.57)	-0.007 (0.43)	0.023 (1.54)	0.012 (0.21) (coefficient for dummy on rule-funds)	0.083 (1.20) (coefficient for dummy on rule-funds)		
Log anticipated funds							-0.102 (2.35)*	0.033 (0.76)
Observations	177	177	176	176	176	176	176	176
R-squared	0.05	0.02	0.15	0.04	0.13	0.03	0.17	0.03
Absolute value of t statistics in parentheses								
* significant at 5%; ** significant at 1%								

Notes: This table shows OLS estimation results for the relationship between unanticipated/anticipated funding and gains in cognitive achievement for English and Mathematics. Columns (1) and (2) report coefficients without any further controls, Columns (3) and (4) include whether the school is rural, whether the head-teacher changed in the previous year, whether the PTA changed in the previous year, differences in PTA fees between 2002 and 2001 and a dummy for private schools. Columns (5) and (6) report results from using a dummy for whether the school received unanticipated funds or not with the same set of controls. Columns (7) and (8) present estimation results when we use the legislated anticipated funds rather than the anticipated funds received at the time of the survey. For the sample, 3 schools with unlikely changes (>2 or <-2 standard deviations) are dropped from sample and the optimal b=3.73 is added on to unanticipated funds to compute the log. All regressions are clustered at the district level. The results from the Wald test for equality of marginal impacts for anticipated and unanticipated funds at different points of the sample range are presented in the text. The results show that we can reject the null hypothesis for English at the 5% level for the sample mean, median, 25^{th} and 75^{th} percentile. We cannot reject the null at these values of the sample for Mathematics.

Table 6: Predicting Discretionary Funds (First Stage)

	(1)	(2)	(3)
	First Stage Hurdle: log unanticipated grants (conditional on receipt)	First Stage Hurdle: log unanticipated grants squared (conditional on receipt)	Probit for Probability of Receipt
Grade VII Male Pass Rate: (Lag 2 years)	-0.021 (1.19)	-0.353 (1.19)	0.017 (2.56)*
Grade VII Female Pass Rate: (Lag 2 years)	0.000 (0.02)	-0.021 (0.07)	-0.014 (2.19)*
Mean School Asset Index (Lagged 1 year)	0.138 (0.23)	2.582 (0.25)	0.358 (1.76)
Is this school eligible for external funds?	-1.606 (1.90)	-27.529 (1.90)	0.885 (3.37)**
Log of District receipts from external sources	1.051 (2.22)*	16.584 (2.05)*	0.327 (1.93)
Log of Province receipts from external sources	0.183 (0.40)	3.463 (0.44)	0.095 (0.56)
Dummy for politically active district	4.581 (2.56)*	89.795 (2.93)**	-0.189 (0.32)
Dummy for politically active district (2)	1.918 (1.24)	43.072 (1.62)	0.943 (1.20)
Constant	0.901 (0.21)	-43.145 (0.60)	-4.455 (2.82)**
Observations	38	38	165
R-squared	0.50	0.52	0.16 (Pseudo R-squared)
Absolute value of t statistics in parentheses			
* significant at 5%; ** significant at 1%			

Notes: This regression shows the first stage of the IV strategy using a hurdle model. Columns (1) and (2) show the estimation results conditional on receipts for unanticipated funding and unanticipated funding squared, while Column (3) estimates the probability of receiving such funding. The log of district receipts from external sources was computed through a questionnaire administered to district authorities and that for provincial receipts through surveys administered at the province level. The two dummies for politically active districts is based on interviews and newspaper articles in the run-up to the election. The predicted value for the second stage is then calculated as E(y)=E(y|receipt) x Prob(receipt).

Table 7: Learning and Funding (IV Results)

	(1)	(2)	(3)	(4)	(5)	(6)
	Hurdle IV: English	Hurdle IV: Math	Hurdle IV: Expected Rule Funds (English)	Hurdle IV: Expected Rule Funds (Math)	Comparison: (OLS, English)	Comparison (OLS, Math)
Hurdle instrumented log unanticipated grants	0.158 (2.71)*	0.090 (2.50)*	0.128 (2.38)*	0.101 (2.91)**	0.082 (2.67)*	0.039 (1.97)
Hurdle Instrumented log unanticipated grants squared	-0.015 (2.27)*	-0.008 (2.45)*	-0.013 (2.28)*	-0.009 (2.81)**	-0.006 (2.30)*	-0.002 (1.03)
Log of anticipated grants	-0.024 (1.12)	0.008 (0.42)	-0.110 (2.43)*	0.038 (0.82)	-0.012 (0.60)	0.024 (1.24)
Observations	164	164	164	164		
R-squared	0.16	0.03	0.18	0.04		

Notes: This table shows second stage IV estimates for the relationship between unanticipated/anticipated funding and gains in cognitive achievement for English and Mathematics. Columns (1) and (2) report coefficients include controls for whether the school is rural, whether the head-teacher changed in the previous year, whether the PTA changed in the previous year, differences in PTA fees between 2002 and 2001 and a dummy for private schools. Columns (3) and (4) report results using legislated anticipated funds rather than the anticipated funds received at the time of the survey. Columns (5) and (6) report results for comparison with OLS results from the same sample. For the sample, 3 schools with unlikely changes (>2 or <-2 standard deviations) are dropped from sample and an additional 13 schools are dropped due to lack of data at the district level. All regressions are clustered at the district level. The results from the Wald test for equality of marginal impacts for anticipated and unanticipated funds at different points of the sample range are presented in the text. The results show that we can reject the null hypothesis for English at the 1% level for the sample mean, median, and 25th percentile and at the 5% level for the 75th and 90th percentile. For Mathematics we can reject the null at the 5% level for the 25th percentile and the 10% level for the sample mean and median. We cannot reject the null for the 75th and 90th percentiles.